INTEL CORE I9-11900K
Team Blue strikes back but is it enough? PG. 76

PHOTOSHOP SKILLS
Amp up your photos with easy SFX PG. 58

BEYOND ENIGMA
The team behind the first CPU PG. 46

MAXIMUM PC

MINIMUM BS • MAY 2021 • www.maximumpc.com

VIRTUAL REALITY

The state of play in 2021

- Tested: The best headsets you can buy
- Oculus versus Vive: Which is the true virtual king?
- A deep dive into VR and game development

STEP-BY-STEP GUIDE
Build this epic **ITX VR-ready** gaming PC!
PG. 68

CRYPTO-MINING BOOM OR BUST!
Find out how to mine your own coins today PG. 26

FUTURE Digital Edition

CORSAIR

STEP UP YOUR GAME

The latest games at their highest settings are no match for the cutting-edge **CORSAIR VENGEANCE i7200 Series Gaming PC.** Powered by a blisteringly fast 10th Gen Intel® Core™ processor and ultra-lifelike NVIDIA® GeForce RTX™ 30 Series graphics, complete with a full range of celebrated CORSAIR components, VENGEANCE PCs take gaming to new heights.

/CORSAIR

VENGEANCE

THE GAMING PC FROM CORSAIR

Find out more at CORSAIR.COM

SAGER
GAMING NOTEBOOKS

SAGER RECOMMENDS WINDOWS

BLAZING PERFORMANCE

$3,349

NP9670M NOTEBOOK
FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- **10th Gen Intel® Core™ i9-10900K Processor**
- 17.3" Thin Bezel Full HD, 144Hz, 72% NTSC Matte Display with G-SYNC
- Opt. 4K UHD Matte Display with G-SYNC
- NVIDIA® GeForce™ RTX 2080 Super 8GB GPU
- Opt. GeForce™ RTX 2070 Super or RTX 2060 GPU
- 32GB DDR4-3200MHz Memory
- 1TB NVMe M.2 SSD
- RGB Per Key illuminated Keybaord with N-Key Rollover
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT
- Sound Blaster ATLAS Audio with Super X-Fi Technology

$1,599

NP8358F2 NOTEBOOK After $150 Instant Savings!

- **Windows 10 Home**
- **10th Gen Intel® Core™ i7-10875H Processor**
- 15.6" Thin Bezel Full HD, 144Hz 72% NTSC Matte Display
- Opt. 17.3' FHD 144Hz 72% NTSC, G-SYNC Display (NP8378F2)
- NVIDIA® GeForce™ RTX 2070 Super 8GB GPU
- 16GB DDR-4 3200MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

$1,499

NP8752D2 NOTEBOOK FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- **10th Gen Intel® Core™ i7-10875H Processor**
- 15.6" Thin Bezel Full HD 144Hz, 72% NTSC Matte Display
- NVIDIA® GeForce™ RTX 2060 6GB GPU
- Opt. GeForce™ RTX 2070 GPU (NP8752F1), RTX 2080 Super GPU (NP8752N2)
- 16GB DDR4-3200MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

All Sager systems come with one year Limited Warranty & Lifetime Tech Support. Prices and specifications are subject to change without notice or obligation. Sager Corporate Offices 18005 Cortney Court, City of Industry, California 91748 ©2018 by Midern Computer, Inc. All rights reserved. Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. All company and/or product names mentioned herein are trademarks and/or registered trademarks of their respective companies. *Free UPS Ground Shipping valid to contiguous US order only.

Windows 10

intel CORE i7 10TH GEN

NP8770F1 NOTEBOOK

FREE UPS GROUND SHIPPING*

$1,699

- Windows® 10 Home
- 10th Gen Intel® Core™ i7-10875H Processor
- 17.3" Thin Bezel Full HD 144Hz, 72% NTSC Matte Display
- NVIDIA® GeForce™ RTX 2070 8GB GPU
- Opt. RTX 2060 GPU (NP8770D2), RTX 2080 Super GPU (NP8770N2)
- 16GB DDR-4 3200MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

NP7877DW-S NOTEBOOK

After $150 Instant Savings!

$1,499

- Windows® 10 Home
- 10th Gen Intel® Core™ i7-10875H Processor
- 17.3" Thin Bezel Full HD 144Hz Matte Display
- NVIDIA® GeForce™ RTX 2060 6GB GPU
- 32GB Dual Channel DDR4-3200MHz Memory
- 1TB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

NP7858DW NOTEBOOK

After $100 Instant Savings!

$1,299

- Windows® 10 Home
- 10th Gen Intel® Core™ i7-10875H Processor
- 15.6" Thin Bezel Full HD 144Hz Matte Display
- NVIDIA® GeForce™ RTX 2060 6GB GPU
- 16GB DDR4-3200MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

NP7958C NOTEBOOK

FREE UPS GROUND SHIPPING*

$1,299

- Windows® 10 Home
- AMD Ryzen™ 7 3700X Processor
- Opt. AMD Ryzen™ 9 Processor
- 15.6 Thin Bezel Full HD 144Hz Matte Display
- NVIDIA® GeForce™ GTX 1660Ti 6GB GPU
- Opt. GeForce™ RTX 2070 GPU (NP7958F1)
- 16GB DDR4-3200MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX200 + BT

NP6877CQ NOTEBOOK

FREE UPS GROUND SHIPPING*

$1,249

- Windows® 10 Home
- 10th Gen Intel® Core™ i7-10750H Processor
- 17.3" Thin Bezel Full HD 144Hz Matte Display
- NVIDIA® GeForce™ GTX 1660Ti 6GB GPU
- Opt. GeForce™ GTX 1650 GPU (NP6877BQ), GTX 1650Ti GPU (NP6877EQ)
- 16GB DDR4-2933MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

NP6858BQ NOTEBOOK

FREE UPS GROUND SHIPPING*

$1,049

- Windows® 10 Home
- 10th Gen Intel® Core™ i7-10750H Processor
- 15.6" Thin Bezel Full HD 144Hz Matte Display
- NVIDIA® GeForce™ GTX 1650 4GB GPU
- Opt. GeForce™ GTX 1650Ti GPU (NP6858EQ), GTX 1660Ti GPU (NP6858CQ)
- 16GB DDR4-2933MHz Memory
- 500GB NVMe M.2 SSD
- Single Zone Multi-Color Backlight Keyboard
- Intel® Dual Band Wireless Wi-Fi 6 AX201 + BT

CALL US TOLL FREE: 1.800.669.1624 **CUSTOMIZE & PRICING GO ONLINE: SAGERNOTEBOOK.COM**

Dealer/Var, government and corporate pricing are available. Please call for details.

SUBSCRIBE TODAY!
see PG. 34

inside

MAY 2021

18 VIRTUAL REALITY

26 THE WORLD OF CRYPTOCURRENCIES
Digital cash, mining, and a grim look at the future of graphics cards.

36 A POTTED HISTORY OF LINUX
Or how a 21-year-old's bedroom coding project took over the world and a few other things on the way.

46 A COLOSSAL SECRET
British WWII braniacs helped save millions of lives while also inventing the computer.

QUICKSTART

10 THE NEWS
Rocket Lake; chip shortage reaches phones; cameras hacked; more.

13 TECH TALK
Intel's Rocket Lake is preparing for launch.

Turn your gaming dreams into virtual reality with a VR-ready rig.

R&D

55 HOW TO
Use Photoshop (part 3); manage Docker containers; crack the Enigma code on a Pi; and more.

56 AUTOPSY
Take a spine-tingling peek at the Backbone One's innards.

68 BUILD IT
A step-by-step guide to building an ITX VR-ready gaming machine.

LETTERS

16 DOCTOR

94 COMMENTS

IN THE LAB

76 INTEL CORE I9-11900K

78 GIGABYTE AORUS 17G

87 CORSAIR K70 RGB TKL

89 RAZER KIYO PRO

90 EVERSPACE 2

© ROCKFISH GAMES

LOOK FOR IT ON DIGITAL,
4K ULTRA HD™
AND BLU-RAY™

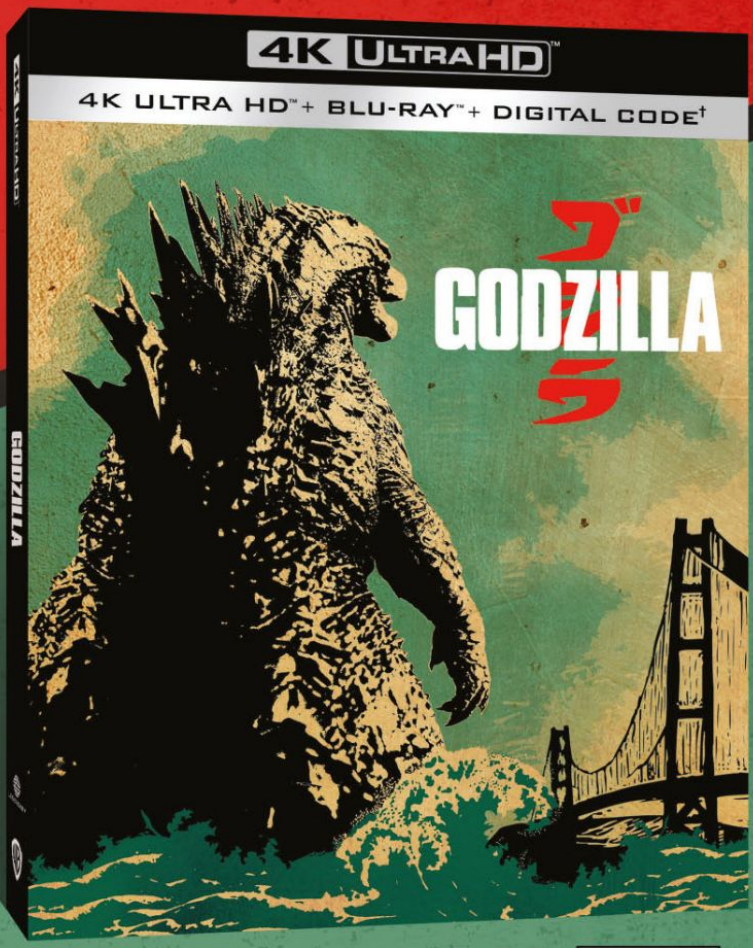

PG-13 INTENSE SEQUENCES OF DESTRUCTION, MAYHEM AND CREATURE VIOLENCE

SCAN TO SEE MORE!

@GodzillaMovie @GodzillaMovie #GODZILLA4K

© 2014 WBEI, Legendary Pictures Productions, LLC and Ratpac-Dune Entertainment LLC. GODZILLA TM & © Toho Co., Ltd.

PUT A PAUSE
IN YOUR DAY
With so many demands from work, home and family, there never seem to be enough hours in the day for you. Why not press pause once in a while, curl up with your favourite magazine and put a little oasis of 'you' in your day.
PRESS PAUSE
ENJOY A MAGAZINE MOMENT
To find out more about Press Pause, visit;
pauseyourday.co.uk

MAXIMUMPC

EDITORIAL
Editor: Zak Storey, zak.storey@futurenet.com
Staff Writers: Christian Guyton, Sam Lewis
Contributing Writers: Dave Alcock, Ian Evenden, John Knight, Jeremy Laird, Chris Lloyd, Neil Mohr, Nick Peers, Brittany Vincent, Jarred Walton
Copy Editor: Katharine Davies
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Photography: Phil Barker, Olly Curtis, Neil Godwin, Dave Alcock
Cover Photo Credits: Intel, Adobe, International Spy Museum

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Mike Peralta mike.peralta@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com
West Coast Account Director: Jack McAuliffe, jack.mcauliffe@futurenet.com

Director, Client Services: Tracy Lam, tracy.lam@futurenet.com

PRODUCTION
Head of Production: Mark Constance
Production Manager: Vivienne Calvert
Project Manager: Clare Scott
Production Assistant: Emily Wood

FUTURE US, INC.
11 West 42nd Street, 15th Floor, New York, NY 10036, USA
www.futureus.com

INTERNATIONAL LICENSING & SYNDICATION
Maximum PC is available for licensing and syndication. To find out more, contact us at licensing@futurenet.com or view our available content at www.futurecontenthub.com.
Head of Print Licensing: Rachel Shaw, licensing@futurenet.com

SUBSCRIBER CUSTOMER SERVICE
Website: www.magazinesdirect.com
Tel: 844-779-2822
New Orders: help@magazinesdirect.com
Customer Service: help@mymagazine.co.uk

BACK ISSUES
Website: https://bit.ly/mpcsingleissue

Next Issue On Sale May 25, 2021

© 2021 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

FUTURE
Connectors. Creators. Experience Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief executive **Zillah Byng-Thorne**
Non-executive chairman **Richard Huntingford**
Chief financial officer **Rachel Addison**
Tel +44 (0)1225 442 244

Zak Storey

VIRTUAL FUTURES

I'VE ALWAYS had a soft spot for virtual reality, but it stems perhaps from something a little more abstract than what we can see today. Back in the day, when I was but a mere toddler, wobbling around my family home, my old man would religiously watch *Star Trek: The Next Generation* on TV every week. I grew up with that and the voice of Jean-Luc Picard, along with other sci-fi series and actors, in the background, and became obsessed myself. The one thing in particular that always fascinated me with almost all of the *Star Trek* sagas was the idea of a Holodeck. A virtual playground in which light itself could be warped into any form and take you anywhere you wanted.

Gene Roddenberry's vision of the far future was equal parts terrifying and exciting, and likely a long way away from where we're headed, but the basic concept of the Holodeck, and by extension virtual reality, is inherently enticing to anyone looking for an escape from the mundane everyday. VR headsets are the first step on that path, and it got us thinking: Although virtual reality exploded with much fanfare back at its launch in 2016, we haven't heard much about it since. So, I set Sam to task with one simple question: Where are we today? And I'm glad to say, the young 'un's done good, with a full eight-page feature on where we are, how we got there, and how the latest offerings stack up to the virtual gods (i.e. me, you, and every other PC enthusiast).

To go alongside that, I got my hands dirty this issue and went on the hunt for some VR greatness myself, but in the form of pure performance, aka a super-small-form-factor ITX build, packing some kick-ass tech, housed in Phanteks' latest and greatest Shift 2 Air chassis. It was a doozy to build, that's for sure, and I hope you stick around for that one!

If that's not your fancy, Jarred's on task this issue with an in-depth look at the state of cryptocurrency, how to mine it, and the ramifications of that for the world at large. Neil gives us a complete tour of the origins of Linux and operating systems as a whole, while Ian introduces us to the men and women who cracked one of the most advanced forms of encryption you could possibly imagine in World War II, by accidentally inventing the first CPU nearly 30 years before Intel did. Honestly, it's an incredible piece that last one, and just goes to show how history has got us where we are today.

Speaking of CPUs, we have a review of Intel's latest and greatest Rocket Lake Core i9-11900K, of course, plus a first look at an RTX 3000-series Gigabyte laptop, and a bevy of other reviews, tutorials, columns, builds, and more, including a little four-page follow-up guide to last month's NAS OS installation feature, concerning Docker, from our good friend Nick Peers.

I do hope you enjoy the latest issue of *Maximum PC*. As always, stay safe out there, and until next time....

Zak Storey is Maximum PC*'s editor and long-time staff member. He's been building PCs since he was 10, and is more than capable of butting heads with the biggest names in tech.*

↘ submit your questions to: comments@maximumpc.com

THE NEWS

Rocket Lake Lands

Intel's desktop chips finally get a new core

IT FEELS as though we've been waiting an age for this. Intel has finally released a desktop CPU with a new microarchitecture: Rocket Lake. It carries a Cypress Cove core, a derivation of Sunny Cove, which was first unveiled in 2018. Originally destined to be a 10nm design, due to myriad technical reasons, Intel had to change tack and "back port" the design to 14nm. So, what we have is the CPU from an Ice Lake chip, combined with Tiger Lake's GPU, all rejigged on to a 14nm die. It's a clumsy compromise. AMD smoothly moved to 7nm nearly two years ago. Officially, the new arrivals are 11th Generation S-Series processors.

Whenever a new CPU microarchitecture arrives, the first question is: What is the IPC (Instructions Per Clock) improvement? Intel claims it is 1.19. This is reasonable, and comparable to the steps AMD's Zen design has made. The next question is: How many cores? This is one of AMD's strong suits, of course. All the new Core i9 and i7 chips have eight cores, while the Core i5 range has six. Intel reckons the IPC lift should compensate for the lower core count in the Core i9s, while it should give a reasonable bump for the Core i7 and i5 chips, where the core count is the same. AMD's 16-core beast is $260 on top of Intel's best, but the 12-core Ryzen 9 5900X is only $10 extra for four more cores. Luckily for Intel, game performance currently tends to stutter after you reach eight cores, but for content creation, it's a real advantage. Why stop at eight cores? Apparently Intel ran out of space on the die, if this is to be believed.

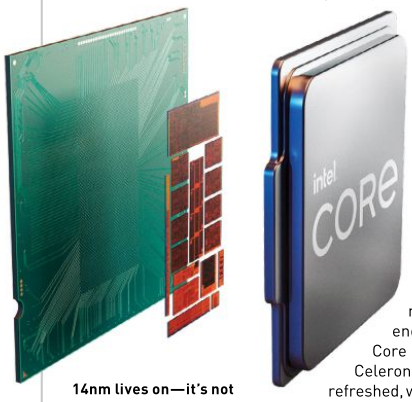

14nm lives on—it's not what Intel planned, but at least we get a version of Sunny Cove on the desktop.

There are 20 or so new versions, all middle and high-end models. The Core i3, Pentium, and Celeron brands are being refreshed, with slightly higher clock rates, but no Cypress Cove cores or other goodies. At the top of the tree is the Core i9-11900K, with a base clock of 3.5GHz, an all-core boost of 4.8GHz, and a maximum single-core boost of 5.3GHz given perfect conditions. Intel has a confusing number of "boost" technologies; it lists four variations for the i9-11900 chips in the promotional material. These include a new Adaptive Boost Technology, which can dynamically boost speeds given enough thermal headroom. There has been little real movement in clock rates, bar a modest 100Hz here and there, so this is one area where Intel does score—sometimes there is no substitute for raw clock speed.

There are three unlocked versions, available with or without integrated graphics. These have a TDP of 125W. The rest of the range is a more modest 65W, apart from the slowest "T" variants, which come in at 35W, thanks to significantly slower clocks. As ever, these TDP numbers are little more than a rough guide—pressure a 65W chip and it'll soon draw much more than that. Rocket Lake also brings us integrated Intel Xe graphics, dubbed Intel UHD 750 or 730, and PCIe 4.0 support. Prices range from $157 to $540. There are some small price increases over the old models. New processors mean new motherboard chips, too. If you have a 400-series board, you may well be able to upgrade; if not, there's a new set of 500-series boards sporting the required LGA 1200 socket, PCIe 4.0 lanes, and more.

Is this a solid response to AMD's Ryzen? Well, it's an answer of sorts, but even hardened Intel adherents won't be cheering too loudly. Intel has kept pace, but not much more. The new microarchitecture is welcome, but without a corresponding shrink, it hasn't delivered the jump Intel planned. As clock speeds and core count haven't moved much, it all boils down to that IPC improvement. It's too early to judge this yet, but not every early benchmark shows the kind of improvement Intel is talking about. Intel still can't match AMD's finest clock for clock. There have also been early worries over thermal performance; these chips can get hot under load. The design is intended for 10nm and there have been compromises.

In today's climate, concerns about relative performance aren't too important in the big picture. Both Intel and AMD can sell what they make. Intel has lots of experience at 14nm, and has its own fabrication plants, which should help with costs. It also means it doesn't have to book production slots at busy fabricators. AMD has to compete with Apple and Nvidia for time at TSMC plants. Intel has done enough to stay king of the single thread, but by most other metrics, it's only holding its own. Rocket Lake chips are in our test rigs now, so you'll have the full story next month. **–CL**

© INTEL

CHIP SHORTAGE REACHES PHONES

ANOTHER INDUSTRY STARTS TO FEEL THE PINCH

PROBLEMS with the worldwide supply of chips have reached Qualcomm. Increased demand, exacerbated by US-China trade relations, and the fact that Samsung's plant in Texas had production disrupted by power outages (which cause problems that take time to fix) mean there are significant shortages across the Qualcomm range, specifically the Snapdragon chips that power smartphones. As prices increase, companies order more to avoid paying more later, and it gets worse. By the second quarter, the shortages will reach the market, with 5G phones hardest hit—production is due to fall by around 30 percent. Elsewhere, shortages continue to worsen. General Motors has even started selling some pickups minus a fuel management chip. Lead times on some minor components have reached 11 months; other vital chips, such as Wi-Fi, Bluetooth, and power management ICs, are five months behind; even CPUs and memory are running at three months. It won't be sorted out before the end of the year, so meanwhile expect higher prices or a longer search—probably both. **–CL**

NEW UPDATE, NEW CRASH

Win 10 manages to nix printing

WINDOWS UPDATES have a reputation for causing somebody, somewhere, a headache. A recent automatic security update (KB5000802 specifically) proved no exception, and the headache was the Blue Screen of Death. The issue revolved around printer drivers for many major brands, including Kyocera and Ricoh. If you tried printing, your PC died, comprehensively. You could uninstall the update, which worked, mostly, but it had to be done individually. Since updates bundle all sorts of things together, including important security updates, this isn't ideal. If this sounds like a small problem, imagine being in charge of a company with hundreds of machines that go bang every time somebody tries to print something. System administrators quickly made their feelings known on forums. Don't worry now, though—yet another update followed a few days later to fix the problem. **–CL**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

GAMES GET HDR
Microsoft is to bring the Xbox's Auto HDR feature to Win 10—about 1,000 games will feel the benefit of deeper, richer colors.

WANT A PC SWITCH?
The crowdfunded Aya-Neo is here. It follows Nintendo's lead, and has a 7-inch screen with AMD silicon—but costs $689.

AMD'S MEMORY BOOST
Smart Memory Access is coming to 3000-series CPUs; expect a 3-10 percent games boost.

TRAGEDIES

APPLE'S LITTLE FAILURE
Production of the iPhone Mini has been cut by 70 percent; it only accounts for about 6 percent of all iPhone sales.

MINERS WANT LAPTOPS
Cryptocurrency miners are buying up gaming laptops, solely for the RTX 3000 GPUs. Bah.

INTEL'S HEAVY FINE
Courts have hit Intel with a $2.18 billion judgment over patent infringements, equal to a month's income. It will go to appeal.

CHEAPEST RDNA 2 CARD YET

Radeon RX 6700 XT joins the family

AMD HAS A NEW RDNA 2 CARD, the RX 6700 XT, which joins the range at the bottom. Essentially a replacement for the 5700 XT, it competes between the Nvidia RTX 3060 Ti and RTX 3070, or so AMD would have it. The GPU has 40 compute units, 160 texture units, 2,560 stream processors, and 64 ROPs. The base clock speed is 2,321MHz, with a maximum boost of 2,581MHz. It carries 12GB of GDDR6 over a 192-bit bus, giving a maximum memory bandwidth of 384GB/s. You also get 96MB of zippy Infinity Cache. This equates to a peak single precision computing power of 13.21TFLOPs. AMD quotes frames rates at 1440p of up to 126fps for *Fortnite*, and 71fps for *Dirt 5* in Ultra High, and it'll chew through 230W at full throttle.

The RX 6700 XT has less memory than its bigger siblings, which all carry 16GB, but it has 4GB more than the Nvidia rivals. While it has two thirds of the compute units of the RX 6800, it runs at significantly faster clock rates. The upshot is a card geared toward 1440p gaming, rather than 4K. If you fire up the ray-tracing option, AMD recommends 1080p. Initial benchmarks look adequate—stick to 1440p, and it'll trade blows with the RTX 3060 Ti, and can match the RTX 3070 at times. The cost is $479—not cheap considering the cut-down chip and relative performance. We'll give you a full review next month.

Nvidia made a fuss about limiting crypto-mining on its newest card, but AMD says it has no plans to limit mining capabilities on its RDNA 2 cards. It claims they are tuned to gaming, and don't make cost-effective mining cards. It has a point, but when there's a quick profit to be had, and a shortage of high-power cards, you can bet 6000-series cards will be worked to death mining digital coins, rather than doing their proper job: entertaining us. **–CL**

© QUALCOMM, MICROSOFT, NVIDIA

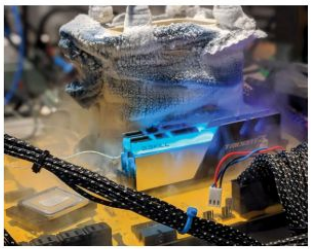

INTEL PULLS OVERCLOCKING PROTECTION

Is the age of overclocking over?

INTEL HAS DECIDED to ends its Performance Tuning Protection Plan. It will honor existing warranties, though, and the Xeon W-3175X will still be covered. Under the PTPP, you could pay an extra $20 or $30, depending on the model, to cover frying your chip while running it outside the published specifications. If things did go pop, you were eligible for a replacement.

Intel claims customers "increasingly overclock with confidence," so there is little demand for PTPP. There is something in that—chips are (nearly) impossible to burn now. More pointedly, the overclocking scene has changed. When PTPP launched, it was the days of Sandy Bridge, and chips were sold with acres of unused headroom. You could take a Core i5-2500K and wring anything up to an extra 1GHz out of it. No modern chip has anything like that level of unrealized performance.

Intel now has strong competition from AMD, and the market for high-end chips is healthy. Virtually every chip has been carefully graded and optimized to reach its maximum potential. This, along with various new boost modes and thermal management systems, means there is often little to be gained by tinkering. Overclocking isn't dead as such—adding some serious cooling remains a viable option—but the days of making a couple of quick tweaks and getting big returns are gone. Mid-range chips can be bumped by 5 to 10 percent for some tasks; the low end has more potential, but that's only useful if budget is the primary concern. Overclocking won't be popular until Intel and AMD sell chips that aren't already running near potential, and in this competitive and speed-hungry market, they don't. **–CL**

150,000 Cameras Hacked

VERKADA OFFERS "hybrid-cloud" security camera systems, and customers include numerous public institutions, as well as some big-name companies, such as Tesla, Nissan, and web security firm Cloudflare. Its server was hacked, and images and video clips totaling 5GB were taken from its 150,000 active cameras and released to the press. These included scenes from inside county jails and hospital.

The breach lasted over 36 hours, and was made by a collective called APT-69420, seemingly simply to make a point. A spokesperson, Tillie Kottmann, cited curiosity, anarchism, and fun as motivations, along with more politicized concerns. The security on Verkada's systems was described as "non-existent and irresponsible." The hackers managed to obtain super admin access, so had free rein to view supposedly secure footage from sensitive locations. Verkada has assured its customers that it is "confident that all customer systems were secured" after the breach went public.

There are a lot of cameras watching us. Legal protection from surveillance is minimal, and based around the concept of an expectation of privacy. This is open to interpretation, leaning toward acceptance. Who has access to footage is similarly poorly defined. Meanwhile, untold amounts of footage of us are piling up on server racks. Now add Internet access to that footage. The hack shows again that going to the cloud without solid security isn't always going to end well. Where we are watched is one debate, but if we are watched, it must be done securely. This hack was done to make a point, other groups may have different concerns. **–CL**

Who's zooming who?

PS5 Winning Console War

Sony's PlayStation 5 is the fastest-selling hardware platform ever in the US, according to analyst NPD. Only Nintendo's Switch sells more. Last year Sony shifted 4.5 million PS5s worldwide, matching the start the PS4 made. The company expects to sell over 100 million in total. Microsoft's Xbox managed around 2.9 million. The entire market is buoyant. Gaming hardware sales reached $406 million in February—the highest since 2011—and games sales hit $4.6 billion, 35 percent more than last year. All this in the face of supply shortages. Microsoft and Sony can sell every console they can make. Part of Sony's success is that it has managed to maintain supplies more effectively. **–CL**

3nm Chips This Year

Despite the troubles in the chip business, and rumors last year, the road to 3nm is still on course. TSMC is to start production this year. This will be "risk" production, where yields are low as problems are addressed. Production will start at 30,000 wafers a month. Full-scale production is expected to start in the second half of next year at 100,000-plus wafers a month. Shifting to 3nm from 5nm brings a transistor density bump of 1.7. Power requirements are an estimated 27 percent lower, and up to an 11 percent speed increase is forecast. The chief recipient of this good news will be Apple, which has booked much of the 3nm capacity. Expect to see the 3nm chip to make its debut in an iPhone. **–CL**

© VERKADA

Jarred Walton

TECH TALK

Rocket Lake Preparing for Launch

I FIRST STARTED WRITING for *Maximum PC* just before Intel launched its Broadwell desktop processors in 2015. Those were the very first 14nm chips to hit the market, and even after delays, Intel was basically two years ahead of the competition. So, it seems practically inconceivable that, nearly six years later, we're now looking forward to the seventh Intel CPU generation to use a variant of 14nm lithography.

Intel stumbled, badly, on the transition to 10nm. Yes, we know 10nm Ice Lake processors have been shipping since 2019, and SuperFIN (10nm+) Tiger Lake CPUs launched last year. But both were confined to mobile solutions, suggesting Intel was still not ready to make bigger 10nm chips. Server-class Ice Lake-LP 10nm Xeon CPUs and desktop and mobile Alder Lake parts will change that later this year. In the interim, we have Rocket Lake. It feels a lot like this will be a repeat of what we saw with Broadwell desktop chips: Too little, too late.

Rocket Lake packs new Cypress Cove CPU cores, an architectural back-porting of Ice Lake's Sunny Cove cores to the 14nm process. Sunny Cove originally used 10nm lithography, and back-porting means Intel had to rework things. There are architectural differences, and a trimming down in CPU core counts compared to existing desktop chips. Comet Lake offers 10-core/20-thread solutions at the top of the stack; Rocket Lake will drop to eight-core/16-thread models. If Intel can achieve the 19 percent higher IPC targets it's aiming for, Rocket Lake stands a chance, but it looks very weak compared to AMD's 12- and 16-core offerings.

The Cypress Cove cores bring upgraded AVX512 support, but that's problematic as it's not widely used. With the right workloads, it can provide a significant boost to performance, but AVX can also cause a massive jump in power use.

Rocket Lake will be Intel's first desktop chip to support PCIe Gen4, too, doubling the bandwidth between the CPU and graphics card. The CPU comes with 20 Gen4 lanes, including four dedicated for the first M.2 slot and 16 for the first PCIe slot, similar to AMD's Ryzen platforms as well as Comet Lake. Intel also doubled the width of the DMI 3.0 connection between the CPU and the PCH chipset, from x4 to x8. That should improve throughput to storage devices and peripherals, but it's only on "select" 500-series mobos—presumably all Z590 boards, but not B560 or H510. 400-series boards are also limited to x4 DMI.

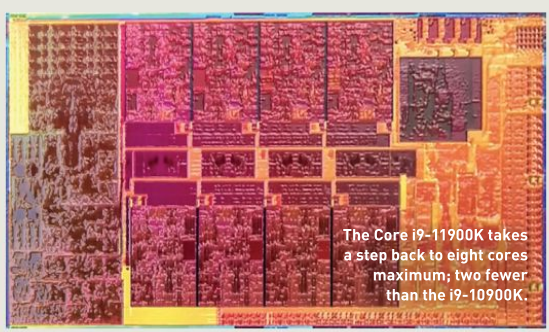

The Core i9-11900K takes a step back to eight cores maximum; two fewer than the i9-10900K.

Finally, Rocket Lake will include Intel's 12th-gen Xe Graphics. All of the Skylake-derived CPUs of the past six years have used some variant of Gen9/9.5 graphics, while Ice Lake laptops included 11th-gen. Tiger Lake moved to 12th-gen Xe Graphics, with up to 96 EUs (Execution Units), and Rocket Lake will also have Xe Graphics, but only up to 32 EUs.

Alder Lake is due this year, though, and its Enhanced SuperFIN 10nm process and Golden Cove cores should beat Rocket Lake. With the 11900K using up to 250W peak power, waiting for Alder Lake seems sensible. The saving grace for Intel is that its 14nm node is mature, with good yields. While AMD looks to increase output of Zen 3 chips, and everyone competes for wafer starts at TSMC, Intel should have plenty of chips to go around.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

© NVIDIA

techradar.pro

IT INSIGHTS FOR BUSINESS

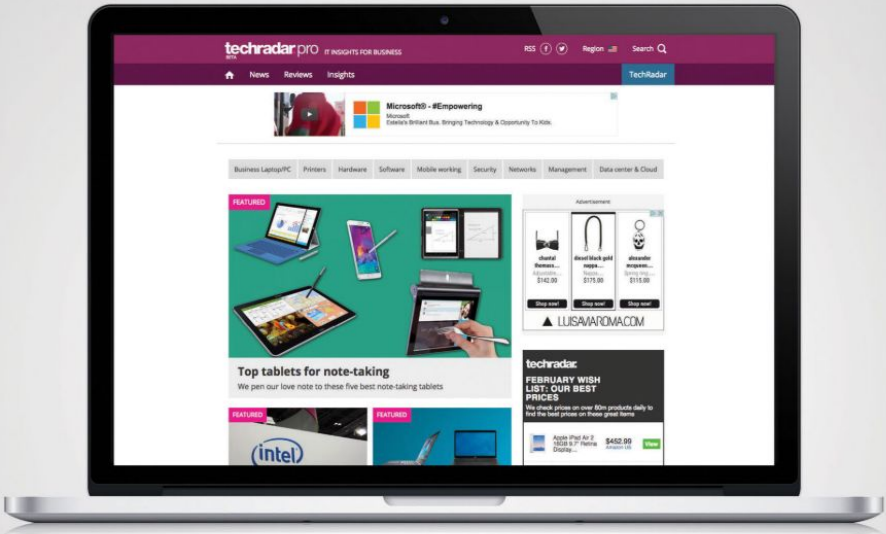

Best practices and solutions

Up-to-the minute news

In-depth reviews

http://www.techradar.com/pro

Brittany Vincent

TRADE CHAT

The Golden Age of PC Gaming

IN LAST MONTH'S TRADE CHAT, I discussed the merits of gaming in a world where consoles are reaching parity with their PC cousins. Long story short, yes—PC gaming is still "worth it" in many ways. But when it comes to playing, archiving, and discovering retro games, it's invaluable.

When it comes to the classics of yesteryear, no platform is doing its part to help preserve and make retro games available more than PC. If you're looking to step into a wonderland of '90s point-and-click adventures or classic shooters, you need only an Internet connection. Simply put, the PC ecosystem is the best possible place to take refuge from the frustrations of modern gaming. If you routinely find yourself saddled with the insatiable desire to dive into the classics, you should regularly be sitting at a mouse and keyboard.

If you're comfortable tinkering and want to spend a few hours at a time doing so, you can always opt to emulate classic operating systems, consoles, and games themselves with a little doing. There are plenty of tutorials (many of them found within the pages of this magazine) that can help you accomplish this feat, but it's not always the end goal for some users. Some players simply want to find a game they enjoyed from 20 years prior and jump in for a fun night of nostalgia.

It's easier than ever to find the games you want to experience from the heyday of PC gaming. And gaming on PC enables you to explore the seemingly never-ending founts of vintage games available across the very marketplaces erected to preserve the titles that brought us where we are in gaming history right now.

Console marketplaces have hosted a modicum of retro titles here and there, but across Xbox Series X, PlayStation 5, and Nintendo Switch, there's a paucity of classic games. While some efforts have been made to offer titles fans want to play from previous console generations, hundreds are missing from these stores. Whether that's due to publishers' refusal to port, licensing issues, or a litany of other reasons, it's an unfortunate reality. You just can't enjoy most retro titles on newer consoles, especially if they came from PC.

> "You can't enjoy most retro titles on newer consoles, especially if they came from PC.

***Nox* is a classic PC dungeon-crawler from 2000.**

What you can do, however, is head to a storefront like GOG.com to browse dozens of classic titles, from the excellent *Phantasmagoria* to lesser-known dungeon crawlers such as *Nox*. It's a well-known repository for games that come wrapped in DRM-free DOSBox containers, so you don't even need to go to great lengths to get these classic titles to run.

Steam is a no-brainer for new titles, but retro games have also slowly made their way to the catalog. Titles such as the previously incredibly rare *Purple Saturn Day* or the difficult to run *The Dame Was Loaded* are now available to purchase for a pittance. Although many collectors (including myself) would prefer hard copies for display, the fact it's so simple to play games that can't be found anywhere else like this is phenomenal. And these two platforms are hardly the only ones to offer such a service.

Consoles are fantastic, and they certainly have their fair share of unique titles. But when it comes to experiencing the golden age of gaming, it'll never be easier to enjoy the oldest, most obscure games on your bucket list than hopping on to your favorite storefront and paying peanuts to experience a blast from the past.

Brittany has been writing about gaming, tech, anime, and more for 14 years. Follow her on Twitter: @MolotovCupcake.

© ELECTRONIC ARTS

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Thunderbolt 3 on PC
> Login Problems
> Moving Folders

Thunderbolt 3 Support

Hi Doc! I've been working from home for about a year. My work laptop is connected to a dual-monitor 180W docking station via Thunderbolt. The dock and all peripherals are mine, but the laptop belongs to the company. I want to use the setup off-hours by unplugging the Thunderbolt from the laptop and into a yet-to-be-built PC rig. How would I do that? I don't need the extra power from the Thunderbolt. Do I just need to make sure the mobo I pick has USB-C? **–John Otto**

THE DOCTOR RESPONDS:
While Thunderbolt 3 is backward compatible with USB-C, USB-C isn't compatible in return. That means you need to source a mobo with Thunderbolt 3 support, which typically means one of two things. First, some higher-end boards come with TB ports on the board—one example is ASRock's X570 Phantom Gaming-ITX/TB3 Mini ITX TB3 ($210, http://amazon.com).

Others provide support for TB3 via an internal header that can be paired with a specific PCIe add-on card, usually from the same manufacturer and paired to specific chipsets—for example, Asus produces two separate cards. The ThunderboltEX 3 ($66, http://newegg.com) card works with older Intel chipsets, while the ThunderboltEX 3-TR card ($101, http://newegg.com) is designed for the newer Z490/B460/H4670/B550 series mobos that have the 14-1 pin header, including the TUF Gaming Z490-Plus ($200, http://bestbuy.com).

Password Problem

I was using my Asus laptop when it did a security update, something it's done perhaps hundreds of times in the past. This time, when the small window for the PIN came up, I got the message, "Something happened and your PIN isn't available. Choose another sign-in option." The only two options it gives me is a password or PIN. I had never set up a password, so what do I do? **–Bruce Noren**

THE DOCTOR RESPONDS:
Windows defaults to using your Microsoft Account credentials to log in, so unless you specifically set up a local user account (harder than it should be) when setting up Windows, you're logging into that account, using your PIN as a shortcut. Try selecting "Password" and entering your Microsoft Account password, which should allow you to log in. If it fails, and you have recently changed your Microsoft Account password for any reason, try the older password instead.

If you don't know what this is, and you don't see an option about "I forgot my password," which will help you try to reset it to gain access, visit https://account.live.com/password/reset in another browser to see if that works. If resetting it doesn't gain you access, you may have to employ the services of a third-party tool that can bypass or reset Microsoft Account passwords in Windows 10. One of the cheapest is PCUnlocker Standard ($19.95, www.top-password.com/reset-windows-password.html).

A select group of motherboards support Thunderbolt 3.

Moving User Folders

Hi Doc. My main drive is an SSD that's getting too full. I installed a secondary HD to hold my music, pictures, videos, and documents. I followed a couple of articles that explained how to point the folders at the top of the "This PC" page ("Music," "Pictures," etc.) to the new drive, but it ended in disaster.

I followed the directions in one of the articles on my "Music" folder, and it successfully moved the folder to the new E drive. The problem is that it deleted all my music files. My big mistake was that I forgot to create a new backup first, so it took me two days to recover everything and get my C drive back to square one, although the association between some of the music files and the cover art was lost.

I don't know if the instructions didn't work because they were too old

© ASUS

↘ submit your questions to: doctor@maximumpc.com

for all the Windows updates that came since, or if it was something I did wrong. Do you know of a simple, foolproof, way to migrate these folders without losing the links in "This PC"? **–Steve Vanetti**

THE DOCTOR RESPONDS: The standard procedure is the same now as it has been since Windows Vista was released, but to avoid the potential for data loss, we suggest you set up the folders on your new drive (for example, "E:\Steve\Documents," "E:\Steve\Pictures," and so on), then copy across the content from your existing user folders to these folders. Once the files have been safely transferred, delete the originals.

Once done, open File Explorer to "This PC." Right-click one of your user folders in the "Folders" list at the top and choose "Properties," then switch to the "Location" tab. Click "Move," browse to and select the target folder on your new drive (so "E:\Steve\Documents" for "Documents"), and click "Apply." You'll be prompted to move existing files to this new folder—this is necessary to ensure Windows understands where the folder resides—so click "Yes," even though you've performed this step already.

Windows is now directed to use the new folder, and you can repeat for any other user folders you wish to move. Going forward, the Doc recommends you put in place a real-time backup of files with versioning support, so you don't encounter this again.

The simplest solution is to tie in these folders to a cloud backup account, but given the likely size of your music folder alone, this may be impractical if you don't have a premium subscription, such as Microsoft 365, which comes with 1TB storage via OneDrive. If you have access to a suitable file server (check out our guide to building your own high-powered NAS from last issue), the perfect tool for the job is NextCloud (https://nextcloud.com). Install the server portion on your NAS—the Doc recommends the Docker version, as outlined on page 48 of the January 2021 issue—and the client on your PC, using it to back up all your key folders in real time. Alternately, go with Syncthing (www.syncthing.net), but be sure to configure the "File Versioning" tab during setup with one of the file-versioning options to keep multiple versions backed up.

CPU Fan Direction

I was puzzled by your answer to Papa Homey's question in January 2021's issue about CPU temperature fluctuation. How did you determine that all his fans were drawing air into the case? I've run Open Hardware Monitor and I don't see any way to tell which way my fans blow. **–Al Kraybill**

THE DOCTOR RESPONDS: If you reread the question, Al, you'll see that Papa himself revealed that all his fans were blowing inward. This obviously isn't something that software can tell you—you need to open the case and visually inspect the fans to see which way they've been positioned to determine if they're pulling air out or drawing it in.

A couple of further points: One, make sure you leave enough space around your PC's chassis to allow the air to circulate, and two, there's no hard and fast rule about whether air should be drawn

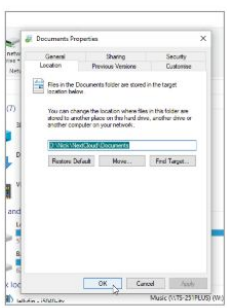

Moved key folders elsewhere? Don't forget to tell Windows.

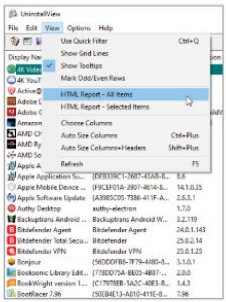

UninstallView generates a complete list of installed apps.

from back to front or vice versa. You may find colder air circulates more naturally at the front, in which case make sure your fans are configured so the air is drawn in from the front and expelled from the rear (don't forget to leave space behind the case). The key thing is to ensure the fans pull colder air across the components to draw the heat from them before venting it.

Installed Software List

My Win 10 install is a couple of years old, so your article (Feb 2021) was welcome. However, the Powershell commands on page 48 didn't work for me, at least not the second one to generate a text file in the root of drive C. I'm not clear about the character(s) following "Format-Table" and end up with a missing argument error. **–William Barnard**

THE DOCTOR RESPONDS: Apologies—there are a couple of errors with the way the command was formatted. It should be one long command, and there should be no space between "-" and "AutoSize." The following is correct:

```
Get-ItemProperty HKLM:\Software\Wow6432Node\Microsoft\Windows\CurrentVersion\Uninstall\* | Select-Object DisplayName, DisplayVersion, Publisher, InstallDate | Format-Table –AutoSize > c:\software.txt
```

If you'd prefer a simpler way that will generate a more comprehensive list of the software on your PC, try NirSoft's portable UninstallView tool, which you can download (choose the 64-bit version), extract, and run from www.nirsoft.net/utils/uninstall_view.html—once opened, a complete list is shown. Choose "View → HTML Report—All Items" to create an HTML document you can refer back to post-installation.

Email Not Working

I can't send or receive emails though Thunderbird anymore. When I check emails via my provider's webmail (https://portal.tds.net), all is well. I've tried contacting TDS tech support by email, and it's no help. All I get are canned responses that are on its site, and that it must be a problem with Thunderbird, as emails can be sent and received via webmail. Any suggestions? **–Rich Taylor**

THE DOCTOR RESPONDS: The problem will likely be with Thunderbird in the sense that TDS may have changed its mail server settings. We found conflicting advice from different TDS websites, but the latest information we could find would seem to point to the server recently having switched to a new system.

Click the "More" button and choose "Account Settings" in Thunderbird. Select "Server Settings" under your TDS mail account. Server Name should be set to mail.tds.net, and the port number will vary according to the Server Type: If it's POP3, this should be set to 995; if it's IMAP, set it to 993. Your username and password should remain unchanged.

Next, select "Outgoing Server (SMTP)," select your existing email settings, and click "Edit." Make sure the Server Name is again set to mail.tds.net, set connection security to STARTTLS (or TLS if you're running an older version of Thunderbird), and verify the port is set to 587. Click "OK" and hopefully you'll be able to use Thunderbird for your email again.

© MICROSOFT

VIRTUAL REALITY IN 2021

How close to a different reality can virtual reality really take us?

Sam Lewis

OVER THE COURSE of the past seven decades, the term "virtual reality" has been thrown around quite a bit within the tech world. However, during that time, it has actually become a rather divisive and subjective form of entertainment. Tedious and awkward are just two ways it has been described. The term itself most certainly encompasses far more now than when the modern take on the technology first debuted in 2016. Heck, today you can get VR headsets for your phone that cost just 10 bucks, while on the flipside you could easily spend over $1,000 getting all the right gear for the ultimate virtual-reality rig.

What VR has been missing since its conception, however, is simplicity and efficiency, and that is seemingly finally starting to change, with better-supported tech and easier-to-use setups.

The bottom line is that there have been many attempts at VR, which have just not taken off and hit the mainstream market in the same way as gaming PCs and consoles have. They aren't as accessible and affordable, and with many setups, you need a powerful PC to run them properly. Space is a massive factor, too, as you need a comfortable area to use it, and this can be a problem for many.

This month, we're taking a look at VR as a whole—where it has come from, some of the latest advances, and the trajectory that this sector of technology is heading in. Will VR headsets finally replace the trusty old monitor? Or is there room for both in the future world we're building?

VR Development

THE NOTION of virtual reality has been used in the tech sector over the last seven decades, but the rate of progress has been slow. Its aim is admirable and enticing: To take you into a fully realized alternate dimension, where you can take a step away from your day-to-day life and become fully immersed in whatever game or media you'd like. When you think about virtual reality, you immediately think of a VR headset/goggles; this concept has been around for years and has stuck true to its original designs in a way.

One of the goals for the technology is that the user should feel as though they are in a different world. This concept could be seen back in the 1950s, with the invention of the Sensorama. Cinematographer Morton Heilig conceived the idea in his 1955 paper "The Cinema of the Future." His design was a machine where the user sat inside a box and watched a movie. This heightened all the user's senses because they were entirely enclosed. By 1962, he had a working prototype that was able to show wide-angle stereoscopic 3D video. Other features included stereo sound, wind, vibrations, and even aromas. The prototype launched with five short films to preview, one of which was a motorcycle ride through New York City. This idea of being able to virtually travel and experience different lifestyles is still partly what fuels the use of virtual reality to this day. Thankfully, *The Walking Dead* wasn't around in the '60s; we can't begin to imagine what aroma they would have used for that.

Skip forward to 1965, and the form factor we see today was starting to take shape. A scientist called Ivan Sutherland started the idea of what we call virtual reality. In his essay titled "The Ultimate Display," he described a computer with a connected display that gave users the chance to gain familiarity with concepts not realizable in the physical world, "a looking glass into a mathematical wonderland." Scientific, we know, but the concept of a "virtual reality" was there.

Moving on to 1968, Sutherland and some of his students created the world's first head-mounted display. Pretty impressive for the time, right? The system was made of two CRT monitors with mirrors and prisms inside. This allowed for digital wireframe graphics to be superimposed over what you could see. The device had six degrees of tracking, with three ultrasonic transmitters on the headset, picked up by four receivers hung from a ceiling. You get the sense that this wasn't a small device by any means. Sutherland named his invention the Sword of Damocles. Unfortunately, it didn't see much progression, probably due to its expense and impracticality.

Nintendo's Virtual Boy made great promises.

The US military's work on vehicle simulation, specialists at NASA, and researchers at the University of Utah were predominantly behind the progress of VR technology in the '60s and '70s. Alongside significant advances in computer graphics technology, VR began to gain heavy traction in the 1980s. By 1985, a team at NASA's Ames Research Center in Mountain View, California, created the Virtual Environment Workstation. This creation was wide-angle stereoscopic, featuring two LCD TV displays and, last but not least, a motorcycle helmet. When life gives you lemons...? The display could be an artificial computer-generated environment or a real environment relayed from remote video cameras. Along with this creation, the term "virtual reality" was conceived by a member of the team that made gloves for the device. The paired gloves could detect movement from the user's fingers via a series of sensors. This created a computer-generated image of your hand in the virtual environment that you could use for interaction. So, even from the mid-1980s, the concept of being able to interact virtually within a simulated alternate reality was already set in stone and would scope the future of VR.

One of the early development kits for the Oculus Rift—business suit not included.

It wasn't until the early '90s that gaming and VR had the opportunity to meet. Arcade machines first brought the experience to the public. Virtuality Group introduced a line of advanced VR machines between 1991 and 1993. Again, they used stereoscopic 3D visuals, with a whopping resolution of 276x372 (and some people complain about VR graphics today). Unfortunately, these arcade systems just didn't gain enough traction. Was the introduction too soon? The performance simply wasn't there to make it a pleasurable enough experience. Sure, it might have been fine for a short arcade blast, but to take it to the next level was a tough ask.

Sega was one of the first companies to attempt to launch a system for use at home. The problem was that the arcade machines were very expensive, so to package the technology up for home use in the '90s was a difficult task, to say the least. Sega VR was announced in 1991—versions of the headset were planned for both arcades and the Sega Genesis console. In 1993, the Sega VR headset was finally announced at CES, but that was the last glimpse of optimism the headset saw. It never materialized, due to complications in production and the fact that it caused severe headaches and motion sickness in its users. Yikes. It was probably for the best that this headset stayed in its prototype form, however, because what happened to the next mainstream VR

release wasn't exactly a great move for this sector of the market.

Nintendo's Virtual Boy was a catastrophic attempt at VR, but it played an important role in the technology's progress. Released in 1995, it aimed to be the world's first portable video game console capable of displaying true 3D graphics. Exciting claims, especially for the mid-'90s. The hype was real on this one. In reality, it was a 32-bit tabletop console that could only show imagery in red and black. Yes, you did read that correctly. It launched at $175 ($302 today, with inflation), which was a lot for a console that couldn't even play in full color. Customers complained about the screen, the lack of software support, and the uncomfortable nature of the headset, which all led to the Virtual Boy's commercial flop.

A NEW ERA IN VR

After that bombshell, VR took a step backward out of the mainstream market, and wasn't to be properly attempted again until 2012. A Kickstarter campaign introduced the Oculus Rift, the first proper progression of virtual reality, stepping away from a gimmicky trend. Palmer Luckey began working on the system between 2009 and 2012. He figured that previous headsets had been very expensive to produce, due to their requirement of complicated lenses that would fix the distortion often associated with having the screen so close to the eyes. He decided to use a cheaper lens but minimize the distortion through clever use of software. Paired with a wider field of view and much better screens, thanks to improvements in technology in the last decade, it created a huge step forward for VR headsets. The Kickstarter raised millions of dollars, clearly showing a huge love for the concept. Valve would help out during this time with its tech knowledge on tracking, to drive PC VR forward. It's true, it did start the re-emergence of VR, and every headset since has followed a similar format. For once, it delivered a promising experience for many.

By 2014, Facebook, interested in the sector, bought Oculus for around $2 billion. Google and Samsung, also attempting a go at VR, developed their own smartphone headsets as well. Google VR consisted of a DIY kit of cardboard and lenses, letting the phone do all the hard work, splitting the image stereoscopically, and using the gyroscope for the tracking. It was a cheap

THE ISSUES

Virtual reality seems to fall short in areas other than what you experience when you're wearing the headset. To be frank, paired with a capable PC and using the latest gear, such as the HTC Vive Cosmos Elite or HP's Reverb G2, which we test in this feature, VR gameplay is very solid and enjoyable on the whole. That, however, is part of the problem. For the vast majority of folk, who aren't as into PCs as us enthusiasts, you also need to get a fairly powerful, usually expensive, rig to pair with your brand-new tethered headset in order to get the premium experience. The high frame rates that give you a smooth VR experience rely on a fairly powerful setup. Running on any old office PC just won't be enough—frame rates will be rather choppy and it will certainly taint your immersive experience. Paired with the already expensive costs of the top-end VR systems, it's simply out of reach for many consumers.

Setting up is another key issue—the process can often be long and tedious, and the tethered systems are not the greatest if you need to take your setup around with you. Also, you're going to want to make sure you have the room for it; you at least need to dedicate a room for the full experience. In many households, this simply isn't a viable option. This, of course, is different from person to person, but VR also is susceptible to issues that many other wearable peripherals encounter.

Wearing a headset for a long duration can be a hot, sweaty, and rather uncomfortable ordeal. Some headsets are pretty heavy, causing strain on the head, neck, and face. It's also easy to get caught in the cables used in tethered systems. And if you are tired or in some form of physical pain, you're not going to be able to get involved in a full VR experience sitting on your couch. So, it can be restrictive in those regards. Nausea and motion sickness have also plagued VR systems forever and are unfortunately unavoidable for some people (especially the editor of this humble publication). If it sounds like it's a fairly long list of negatives, well, it is. So, if you haven't experienced VR before, please try it out before investing in your own setup.

© GETTY IMAGES/ RAWPIXEL

Easily test if your PC can handle VR

and easy way for the public to experience immersion on the go. Nevertheless, the experience was lackluster compared to PC-powered systems, but it was the interest in VR that the phone headsets were capitalizing on that brought more consumers to the market.

2016 was when the market really started to become saturated, and when it gained a solid position in the gaming scene. The companies at the forefront of VR gaming were Oculus, HTC Vive, and Sony, with its PlayStation VR. All of the headsets used motion-tracking controllers for an immersive in-game experience, and for the ability to interact with in-game items, much like the concepts pushed by NASA in the 1980s. These were not stand-alone systems, though—they required a PC or console to deliver the games to the headset. With that came tedious cables and an annoying setup process, but the gameplay was strong. Such systems are known as tethered.

Fast-forward to 2018, and we start to see a more advanced system with the Oculus Go, a stand-alone, seamless-to-use, fully wireless console. It offered a full, more immersive experience, due to its wireless capabilities—but this also resulted in poor battery life. Its library was a little lacking, too, due to its three degrees of freedom capabilities. By 2019, every company's version of their headsets had been updated multiple times, with clearer displays, advanced tracking, and more powerful components. Oculus released the Quest, followed by the Quest 2 shortly after—two more stand-alone wireless consoles, with the latter also being able to pair with a PC. This is definitely the direction the industry needs to head. The ultimate VR experience has to be completely wireless fully powered PC gaming on a headset. Currently, tethered systems deliver the most full-fledged package, but lack the seamless nature of wireless technology.

HP vs. Vive: Com

HOW MUCH DOES it really cost to experience VR with these headsets? We were lucky enough to get our hands on both the HP Reverb G2 and Vive Cosmos Elite headsets. Both of which are great options for solid and seamless VR immersion. With the Cosmos Elite coming in at $300 more than the Reverb G2, which is priced at $599, is the price gap worth it, and how does it affect the overall cost of a setup that also includes a capable PC?

HEADSET TO HEADSET

HP has arrived in the VR market a little later than some of its rivals. Partnering up with Microsoft and Valve, the company's second attempt at a Reverb headset is impressive, to say the least. In many aspects, it stands up to the better-known VR companies, announcing: "I can do that, too." Compared to the Vive Cosmos Elite, it certainly holds its ground. The Reverb G2 doesn't come with any additional tracking—it's all done from the headset itself. This, of course, has its pros and cons. It has the bonus of taking up fewer power outlets, but the tracking on the Vive headset is definitely superior because of this. The overall quality of the G2 isn't as durable and feels a little cheaper than the Cosmos Elite, too, but even this has its perks. A lighter headset means that you can wear it for longer periods. It also

DEVELOPMENT PROBLEMS

We've detailed the hurdles VR needs to overcome to become truly cemented in our technological lives, but there's another angle that's just as important: development—in particular, investment.

Because of the steady speed of sales in VR headsets, game development has certainly come along, but it hasn't seen quite the astronomical leaps the main PC and console market has seen over the last few years. This means we're left with a much narrower market for VR games than standard releases. This further damages the ecosystem of virtual reality, as it's harder for it to gain momentum, because there's no must-have killer app, outside of *Half-Life: Alyx*. The problem is market share, and it's a chicken-and-egg scenario. Because the VR market is rather small compared to the combined might of PC gaming and consoles, it's difficult to attract investors. Why develop a game for a niche system with very specific requirements, which can't be ported to other devices, when you could create a game that can be run traditionally, across PCs and consoles, for the same money, but with a far larger market to sell to?

As a consequence, why buy a VR headset when the game library for it is so limited? And there's the problem: VR needs good games and killer apps to sell more headsets, and killer apps and game development need more headsets sold before they get proper investment.

© THISISENGINEERING RAENG ON UNSPLASH; STORE.STEAMPOWERED.COM

peting Systems Tested and Costed

makes it a little more comfortable than the competition, so it gets a thumbs-up from us.

The two systems also share very similar feature sets. Both have stereo sound through open-back headphones, which hang from the headset. Both are tethered via a dongle that goes into a DisplayPort and a USB C slot on your PC. And both require Windows 10, with support for SteamVR and Windows Mixed Reality as standard.

The controllers, however, do set the two competitors apart. On the Cosmos Elite, the ergonomics feel much more intuitive in VR than with the HP system. The Vive's are comfier to hold, feel better built, and are covered in grippy rubber. The more natural the controller feels, the more immersive the experience. When using HP's controllers, we found them to be a lot less grippy. They're also made from cheaper materials and, sadly, require separate batteries. Come on, guys, it's 2021. That's not to say they are bad controllers—they are light and comfortable—just not up to scratch when compared to Vive's.

What are the actual headsets like to use, though? Well, very comfy, actually. And both are impressive when it comes to clarity, too. We got the best experience with the HP Reverb G2, even though it's the cheaper of the two devices. It has the same 90Hz refresh rate, but the quality of the screen looks better. The resolution is higher and, with slightly smaller screens at 2.89 inches, the G2 takes the crown for pixel density as well. That's not to say the screens on the Vive fall far behind, they're just topped by HP's offering.

Overall, both VR headsets offer a great all-around package, so you won't be disappointed with either option. As you need to take into consideration the cost of each headset, plus what you may need to buy in addition in order to get the best experience, we've come up with a rough price guide on how much everything should cost.

COST ANALYSIS

We virtually built two PCs using PCPartPicker's online tool (https://pcpartpicker.com). We chose parts that would run at the minimum requirements, to give an insight into how much a full build including headset would cost.

The minimum specs for the HP Reverb G2 are an Intel Core i5, 16GB RAM, and at least an Nvidia GeForce GTX 1080 or similar to run at full resolution. Paired with the cheapest components we could find, the whole system comes to $2,460. But bear in mind that graphics cards are all but impossible to buy right now.

We applied the same method for a system that would run the Vive Cosmos Elite. Based on the Vive's minimum requirements of an Intel i5, along with 16GB RAM, and an Nvidia GeForce GTX 970 4GB, we punched the whole lot into PCPartPicker to build a suitable PC. All together, including the cost of the Cosmos Elite itself, this setup comes to $2,600.

Please note that these are merely guides, because pricing for PC parts, especially graphics cards, is constantly changing. More modern cards (such as the GTX 1660 Ti or RTX 3060) should provide far superior performance for far lower cost outside of the hardware drought.

The thing to take away from this is that you're probably going to be spending around $1,500-plus when it comes to a complete VR setup, if you don't already have a solid system. Running on the bare minimum isn't recommended, as it won't give you the immersion you're after. Other costs to consider are the prices of apps and software, or any VR-related subscriptions you may wish to sign up for.

If you want to see how your existing setup stands and whether it's capable of running VR, download the SteamVR Performance Test (http://bit.ly/3cfQHII). Also, see our dedicated Vive Cosmos Elite review (pg. 84) for a more detailed look at one of VR's leading headsets.

Right: **The Vive Cosmos Elite, still one of the best since its debut in 2016.**
Below: **The HP Reverb G2, light, comfortable, and certainly capable.**

10 Great Introductory VR Experiences

SO, YOU'VE BEEN looking into virtual reality and want to know how to get into it, what's on the market, or what could entice you into this new dimension. Because VR is a fairly new technology compared to regular monitors, the market is nowhere near as deep as the traditional gaming one, of course. However, VR is used in other sectors of technology, too. As much as we love gaming, VR can be used for educational purposes, medical advancements, virtual tours, and experiences exploring fashion, and so on—it can be a great form of escapism. Of course, it's exciting to head straight for your intended titles, but sometimes it pays to have a look around at other alternatives or new options, so we've come up with 10 apps and games that will give you a varied VR experience when starting out.

All of these apps and games can help new or existing VR users experience something new and worthwhile to improve their virtual-reality immersion. With the progression of VR support, developers are hopefully going to start creating more rich and immersive content to help grow the industry, because the technology is certainly there to handle such projects.

BEAT SABER

This one was bound to feature in our round-up. We know it's been plastered everywhere, but *Beat Saber* is just so darn addictive. Rather like a futuristic electric twist on *Guitar Hero* (and we mean no offense to the latter), *Beat Saber* is very simple in concept: Slice blocks in two following the arrows shown. But when paired with music and a lot of action, it combines a little fitness and entertainment to get your day started.

• *$29.99, www.steam.com*

HALF-LIFE: ALYX

This one is for the gamers out there. Valve took no shortcuts when it came to bringing the world of *Half-Life* into the VR limelight. It won VR Game of the Year in 2020, so it's obviously something special. Its super-rich immersion makes up hopeful that more games of this caliber will emerge. Plenty of shooting, a very detailed environment, and loads of interaction—what's not to like? This game certainly gets us excited.

• *$59.99, www.steam.com*

THE LAB

This free-to-play experience is a collection of fun mini-games from Valve—it sure has a good foothold in the VR market. It's quite varied, with fun and playful aspects, and also involves some realistic physics, with vivid graphics. The games include *Slingshot, Postcards, Longbow,* and more. One of our favorite activities is Solar System, in which you get to jump into space and experience educational space fun. This app is a good introduction to different aspects of VR and is well worth checking out.

• *$Free, www.steam.com*

TILT BRUSH BY GOOGLE

In Tilt Brush, the world is your palette. This clever app enables you to paint in 3D, which if you haven't tried before, is a very entertaining way to release your artistic side, without all the practical issues involved with more traditional artistic pursuits. You can paint with 3D brush strokes, light, and fire. Having your workspace float around you really takes you away from any external distractions, enabling you to completely focus on your art. It's a great app for experimenting with your creativity.

• *$19.99, www.steam.com*

WEVR BLU

Ever wanted to plunge into the depths of the ocean without any of the many consequences you might encounter in the real world? Well, Wevr Blu is a great way to experience what it's like down there. With this simulation-like viewing, you can get face to face with some of the planet's most amazing species. You can watch a whale swim past you or even get right down to the depths and experience the alien fish that lurk down there. It's an engrossing way to see the potential of VR and what it has to offer.

• *$Free, www.microsoft.com*

©HTTPS://STORE.STEAMPOWERED.COM

JOB SIMULATOR

This is a tongue-in-cheek way to get used to the physics of VR. Its vibrant cartoon-like world is one in which robots have replaced all human jobs, and you experience what it used to be like to work. Jobs include office worker, store clerk, mechanic, and gourmet chef. It's great for getting to grips with the headset and controllers. It's also perfect for any younger family members who want to try VR for fun or educational purposes.

- ***$19.99, www.steam.com***

GOOGLE EARTH VR

Everyone knows what a great tool Google Earth is. It's a perfect way to plan trips, explore your surroundings, or virtually travel to seek out new adventures anywhere in the world. Of course, if you pair this incredible app with virtual reality, you have a much more immersive experience. You can fly over any city you choose, travel to the highest peaks on Earth, and even venture out into space. Sounds good to us.

- ***$Free, www.steam.com***

ARKIO: EARLY ACCESS SOFTWARE

An architect's playground for VR, Arkio combines users to collaboratively create buildings and interiors, and sketch urban plans. You can start fresh or import 3D models. You also get to export your design to other 3D design tools, including Revit, as native Revit geometry. For creatives, this an entertaining way to bring your flare to a fun architectural tool. It's a similar concept to SketchUp, but adds 3D to fully realize your designs virtually.

- ***$Free, www.steam.com***

NVIDIA VR FUNHOUSE

Another great collection of mini-games to bring some fun to your VR immersion. It's super-simple to play and a great throwback to all your favorite carnival games in one small package. Intuitive and graphically impressive, the games include a shooting gallery, balloon popper, flaming archery, basketball, and more. The games are quick and exciting, you complete them in order, and you can view your scores all at the end—a good group game to pass the VR headset around with.

- ***$Free, www.steam.com***

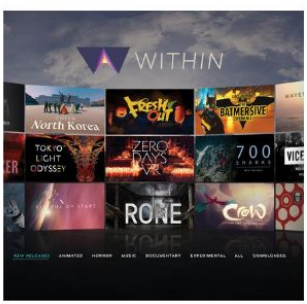

WITHIN VR

For those who aren't particularly big fans of playing videogames, another impressive use of virtual reality is watching 3D content. It adds a whole new level of enjoyment to content, helping you to feel as though you are genuinely there, experiencing scenarios for real. Within enables you to explore and participate in a vast library of stories, award-winning VR documentaries, animation, music videos, horror, and more. It's a great hub for high-quality VR goodness.

- ***$Free, www.steam.com***

HOLODECKS?

AMD said it was looking at the idea of a holodeck in 2013, and start-ups are now joining in, including San Jose's Light Field Lab, which is collaborating with OTOY to develop a display-based holographic experience, backed by Gene Roddenberry's son.

But it's not just private companies looking at this. Turkey's Bilkent University applied some of its finest minds to the problem, and in a paper published on Nature.com (http://bit.ly/MPCHolo) discussed how the poor resolution and fuzziness of current holographic images can be addressed with a better understanding of the math involved. When you have so many vectors and dimensions you need to render at any one time, they cause interference with one another. Professor Ömer Ilday suggests introducing a random wavefront phase to each image to expand the options available, and instead of rendering from three or four projections, you can now use over a thousand to achieve far superior quality without interference.

© HTTPS://STORE.STEAMPOWERED.COM

EVERY FEW YEARS, cryptocurrency returns to the spotlight with a surge in prices, and everyone suddenly wants to invest. This time, anyone with even a passing interest in computer hardware should know what Bitcoin is at some level, but how does it actually work, and when will we be able to buy graphics cards at reasonable prices again?

The first most people heard about Bitcoin was in 2011 when it jumped from being a weird experiment to having actual real-world value, reaching over $30 per BTC. Then it collapsed to below $2 and many were quick to call it a Ponzi scheme. Now, early 2021 has set new record highs, topping out at $58,000.

How far will it fall this time? Will the talk of investments from Elon Musk, Mastercard, and others push it to even higher highs? And why should billionaires even want to bother with cryptocurrencies? Let's find out.

Digital cash, mining, and a grim look at the future of graphics cards. *By Jarred Walton*

THE WORLD OF CRYPTOCU

RRENCIES

© ANDREY RUDAKOV/BLOOMBERG VIA GETTY IMAGES

MEET THE BLOCKCHAIN

The starting point for Bitcoin and other cryptocurrencies is the blockchain. It's become a buzzword, but the short description is pretty drab. The blockchain contains an open ledger of all transactions that have ever taken place on a given cryptocurrency network. Did you send money to Uncle Theo, paying him for a graphics card? It's in the blockchain, for everyone to see. So, when your friend comes by claiming you never paid him back for the concert tickets, or the fast-food stop, or whatever, you can find incontrovertible proof—one way or the other.

That might sound interesting, but it probably sounds alarming, too. Why would you want everyone to be able to check on what amounts to your digital bank account? That's where the second part of the blockchain comes into play: All transactions are at least somewhat anonymous. Unless you publicly list your Bitcoin wallet address—and it's trivial to create more addresses, so you can have more than one—determining what money belongs to whom is difficult. It's not impossible, especially with enough time, effort, and use of a specific wallet, but there are ways to obfuscate your transactions and keep them veiled from prying eyes.

The more important aspect of the blockchain isn't just that it keeps a record of all transactions. The critical factor is that it's backed by strong cryptography—SHA256, in the case of Bitcoin, and dozens of other cryptographic algorithms for other coins. A lot of complex math goes into the creation of a good cryptographic algorithm, with the end goal being that the only way to crack it would be brute force. SHA256 has a 256-bit cypher length, so breaking the encryption for a single message would involve finding the one number out of a 2^{256} number space that's correct.

If that doesn't seem like a big deal, let's put it differently: 2^{256} is the same as 1.158×10^{77}, or in other words, it's a 1 followed by 77 zeroes. Our solar system has roughly 1.2×10^{56} atoms, the Milky Way galaxy has approximately 2.4×10^{67} atoms, and depending on whose math you want to use, the entire universe has anywhere from 10^{78}–10^{82} atoms. That means a 256-bit number space is large enough that you can easily assign a unique number to every atom in the galaxy, with room to spare.

We're simplifying quite a bit, but the main point is that cracking the cryptography behind Bitcoin would be extremely difficult, and other coins use even stronger cryptography. Bitcoin and other cryptocurrencies also use cryptography for the unique user wallet addresses. In the case of Bitcoin, there are currently a potential 2^{160} addresses, or around a quindecillion. No, we didn't just make that word up. That's a one with 48 zeroes after it. Or if you prefer, it's 1 trillion × 1 trillion × 1 trillion × 1 trillion. We're not likely to run out of wallet addresses anytime soon.

All of this cryptography plays a pivotal role in cryptocurrencies, as you might guess considering the name. They're "currencies"—perhaps digital commodities would be a better name, but that ship has sailed—backed by cryptography. And when we say cryptography, what we mean is complex math, the sort of thing to make even your calculus teacher squinch up their face. So-called "fiat" currencies, such as the US dollar, are backed by trust in the government and banking institutions. Bitcoin and Ethereum are backed by our trust in math. Actually, they're backed by math plus all of the people willing to participate in securing the cryptocurrency network, which brings us to the next topic.

SECURING THE BLOCKCHAIN

We've talked about cryptography, but what exactly makes this whole blockchain thing secure? Cryptocurrencies run as a large distributed network of computing resources, with the various nodes communicating with each other. Every time someone wants to move coins—or really, a portion of a coin, with most cryptocurrencies able to deal with amounts to eight decimal places—they send a transaction out on to the network and it becomes available for inclusion in the next block. Each block encapsulates all of the transactions since the previous block into one chunk, creates a digital hash signature of the block, and the entire thing gets submitted as a new addition to the blockchain.

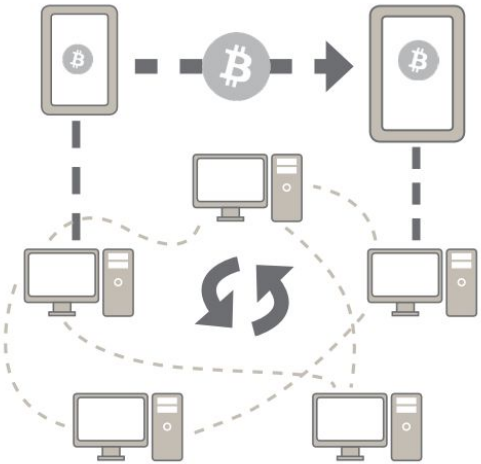

The blockchain is essentially an open ledger keeping track of every cryptocurrency transaction, run over a huge distributed network.

Previous pages: **The BitRiver cryptocurrency mining farm in Bratsk, Russia, with mining rigs packed on to three stories of racks.**

Here's where things get interesting. You can't just add a new block to the network any time you want. That block needs to satisfy certain requirements. In short, you need to find a solution to a mathematical function defined by the coin algorithm. The function uses the results of the previous block as input, going back in a chain to the original genesis block. This means no one can get a head start by trying to solve for future blocks until the current block has been "solved." The algorithms are also designed such that the solutions are effectively random, selected from the entire number space for the cryptography key. There can be a huge number of acceptable solutions, and your hardware just needs to find one of these.

When a system does find a solution, it gets to pick which transactions get included in that block (ideally, all of them, but more on that in a moment), wraps them all up with a tidy bow, calculates a digital signature, and submits that to the network. Other nodes on the network can then verify the solution and accept or reject it. Each potential solution has a timestamp, and the network has various criteria to help determine which block was found first. Eventually, the network reaches a consensus among enough nodes that things move on, so if two different people (machines) find different solutions at the same time, only one ends up as the accepted "winner."

WHAT IS CRYPTOCURRENCY MINING?

Coin mining tends to make a lot of gamers and PC enthusiasts angry—angry or rich, depending on when they started doing it. When you hear talk of GPU shortages and mining farms, it stems from the above cryptography and securing of the networks. At their core, most cryptocurrencies depend on people joining the

SAFE GPU MINING SETTINGS

Assuming you're interested in trying out cryptocurrency mining, there's a lot of conflicting advice on how to go about tuning for optimal performance. Many tutorials throw caution to the wind and look to maximize hash rates, with little heed to the differences in individual graphics cards. It's a fast track to dead hardware if you're not careful. While it's true that AMD and Nvidia graphics cards tend to be tuned somewhat conservatively, the intent is for the card to handle hours of gaming, every day, for several years. 24/7 cryptocurrency mining is a completely different use case.

There's a reason Nvidia's data center GPUs are usually clocked lower than consumer GeForce parts, with very different cooling configurations. Supercomputers use high-speed fans that push a lot of air through the coolers, and the fans are designed to be hot-swappable. The cards are designed for 24/7 high utilization, which means thermal pads and other elements are also built to those specifications. Even someone playing games eight hours a day, every day, won't come close to punishing a card like mining will. Keep that in mind.

Also remember that what works well on one GPU, or even one particular card using a specific GPU, may not work everywhere. The silicon lottery is real, parts are binned and tested to meet certain criteria, and some have more headroom than others. Redlining a graphics card to improve hash rates might result in short-term gains, but it's also a quick way to burn out fans, VRMs, memory chips, or even the GPU. Start slow, check how your card is doing on temperatures and fan speeds, and don't worry about chasing down every last hash.

To check your temperatures and fan speeds, we recommend using HWInfo64 (*image right*). The latest version shows maximum GDDR6X temps on the RTX 3080 and 3090, as well as GDDR6 temperatures on most AMD RX 5000 and RX 6000 cards. (So far, GDDR6 thermals on RTX 20-series and lower tier RTX 30-series GPUs aren't available.) The specific card and fan design play a role, and consumer graphics card fans are not designed to run at a constant 80–100 percent fan speed, all day, every day. Many cards actually limit the maximum fan speed to around 50 percent, and won't go above that unless things get really bad—like super-hot GDDR6X temperatures. Consumer cards also tend to prioritize lower noise levels and higher temperatures, which is counterproductive to 24/7 mining.

If you want a reasonable estimate of where a card should run its fans, turn off any overclock and launch a demanding game such as *Metro Exodus* at 1440p ultra settings, and just let it run for 15–20 minutes, then check temperatures, fan speeds, clocks, and so on. The maximum fan speed you see is what the manufacturer thinks should allow the card to last three-plus years. Go anywhere above that and you're more likely to have fans or other components fail.

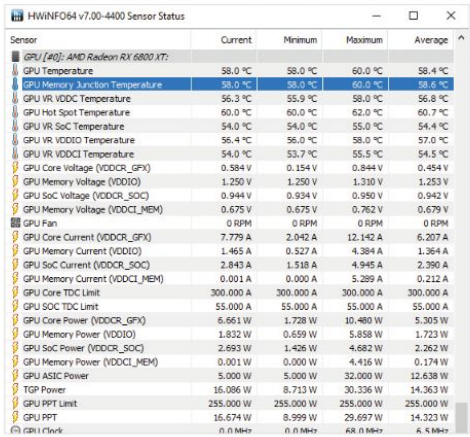

Sensor	Current	Minimum	Maximum	Average
GPU [#0]: AMD Radeon RX 6800 XT:				
GPU Temperature	58.0 °C	58.0 °C	60.0 °C	58.4 °C
GPU Memory Junction Temperature	58.0 °C	58.0 °C	60.0 °C	58.6 °C
GPU VR VDDC Temperature	56.3 °C	55.9 °C	58.0 °C	56.8 °C
GPU Hot Spot Temperature	60.0 °C	60.0 °C	62.0 °C	60.7 °C
GPU VR SoC Temperature	54.0 °C	54.0 °C	55.0 °C	54.4 °C
GPU VR VDDIO Temperature	56.4 °C	56.0 °C	58.0 °C	57.0 °C
GPU VR VDDCI Temperature	54.0 °C	53.7 °C	55.5 °C	54.5 °C
GPU Core Voltage (VDDCR_GFX)	0.584 V	0.154 V	0.844 V	0.454 V
GPU Memory Voltage (VDDIO)	1.250 V	1.250 V	1.310 V	1.253 V
GPU SoC Voltage (VDDCR_SOC)	0.944 V	0.934 V	0.950 V	0.942 V
GPU Memory Voltage (VDDCI_MEM)	0.675 V	0.675 V	0.762 V	0.679 V
GPU Fan	0 RPM	0 RPM	0 RPM	0 RPM
GPU Core Current (VDDCR_GFX)	7.779 A	2.042 A	12.142 A	6.207 A
GPU Memory Current (VDDIO)	1.465 A	0.527 A	4.384 A	1.364 A
GPU SoC Current (VDDCR_SOC)	2.843 A	1.518 A	4.945 A	2.390 A
GPU Memory Current (VDDCI_MEM)	0.001 A	0.000 A	5.289 A	0.212 A
GPU Core TDC Limit	300.000 A	300.000 A	300.000 A	300.000 A
GPU SOC TDC Limit	55.000 A	55.000 A	55.000 A	55.000 A
GPU Core Power (VDDCR_GFX)	6.661 W	1.728 W	10.480 W	5.305 W
GPU Memory Power (VDDIO)	1.832 W	0.659 W	5.858 W	1.723 W
GPU SoC Power (VDDCR_SOC)	2.693 W	1.426 W	4.682 W	2.262 W
GPU Memory Power (VDDCI_MEM)	0.001 W	0.000 W	4.416 W	0.174 W
GPU ASIC Power	5.000 W	5.000 W	32.000 W	12.638 W
TGP Power	16.086 W	8.713 W	30.336 W	14.363 W
GPU PPT Limit	255.000 W	255.000 W	255.000 W	255.000 W
GPU PPT	16.674 W	8.999 W	29.697 W	14.323 W
GPU Clock	0.0 MHz	0.0 MHz	68.0 MHz	6.5 MHz

For temperatures, lower is better, at least up to a point (i.e. not freezing). For long-term mining, we'd try to keep GPU temperatures below 70 C, but that may not even be the critical factor. Keeping VRM (Voltage Regulator Modules) temperatures below 85 C is also desirable, and we definitely wouldn't run with memory temperatures over 100 C and expect a card to remain viable. The RTX 3080 and 3090 are the worst culprits, with GDDR6X often settling in at 106–110 C, but better cards—or ones where you've replaced the default thermal pads with something better—can keep GDDR6X mining temperatures under 90 C.

Depending on your GPU, tuning performance can often dramatically cut power requirements without impacting hash rates, and lower power generally means lower temperatures as well. Ethereum mining prefers memory bandwidth over raw computational power, so turning down GPU clocks and power limits while overclocking the memory usually leads to better performance. Undervolting the GPU core can also help reduce power and heat, particularly on certain AMD GPUs. Search the web for details on your specific GPU, but again be conservative, and don't jump straight to maximum performance settings. If your particular card ends up 10– 15 percent slower on hash rates than you see others reporting, but your temperatures and power use are lower, the card will likely last longer.

From left to right: **Ethereum, launched in 2015; Bitcoin Cash, which split from Bitcoin in 2017; Litecoin, launched in 2011, making it an old-timer; Dash, which is similar to Bitcoin, but uses a two-tier structure of miners and masternodes; Bytecoin—not to be confused with Bitcoin.**

© WIKIMEDIA, HWINFO64

WHAT ARE ASICs?

We've focused most of our efforts on the description of mining with GPUs, but originally Bitcoin was designed to be mined using CPUs. It just so happened that with GPUs becoming increasingly programmable, software developers figured out they could get 10, 100, or maybe even 1,000 times the performance of a CPU by doing the calculations on a GPU.

But GPUs aren't optimized for doing just cryptographic hashing calculations—there's a whole bunch of extra stuff for textures, 3D graphics, anti-aliasing, ray tracing, and so on. What if you stripped out all of those unnecessary capabilities and focused on creating a chip that was only good at one thing: cryptographic hashing? It would be faster than a GPU, potentially mining even more Bitcoins! Not surprisingly, lots of others felt like this was a good idea as well.

The result is called an ASIC, Application Specific Integrated Circuit. Technically, a GPU is a type of ASIC, it just happens to be focused on graphics work, while a CPU is a general-purpose ASIC for x86 architectures, or ARM, or whatever. Bitcoin ASICs focus on boosting SHA256 hashing rates as high as possible. Current ASICs can perform somewhere around 1,000 times more hashes per second than even the best GPUs, using a similar amount of power.

There was also an intermediate step between GPUs and ASICs: the FPGA (Field Programmable Gate Array). FPGAs are reconfigurable silicon that can be updated to run different algorithms. There was a brief span of time when FPGAs were used for Bitcoin mining, offering about 10 times the performance of a GPU in the same power envelope, but it wasn't long before fully custom ASICs replaced FPGAs for most hashing algorithms.

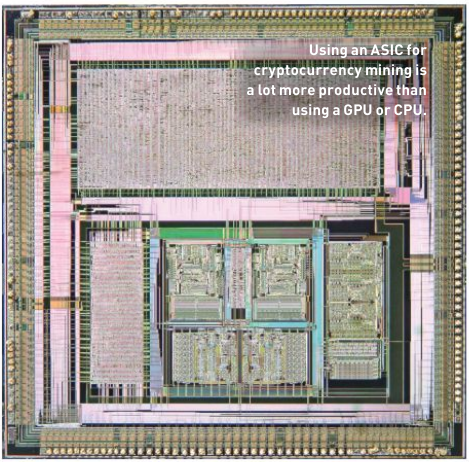

Using an ASIC for cryptocurrency mining is a lot more productive than using a GPU or CPU.

coin network to provide security. What they're actually doing is donating hardware cycles to run the cryptographic algorithms to find the block solutions that keep everything flowing. But why would anyone want to do this? Answer: They get rewarded.

As an incentive to donate your hardware cycles to a particular network, whenever a new block solution is found and accepted, the algorithm provides a block reward. For Bitcoin, this started at 50 BTC per block, while Ethereum provided 5 ETH per block. But there's more to it than just finding a block and getting a reward.

To constrain supply, most coins also reduce the block rewards over time. Bitcoin's algorithm cuts the reward in half every 210,000 blocks—we'll get to the why of that figure in a moment. Dogecoin had a random reward of 0–1,000,000 DOGE per block at first, cutting the maximum reward in half every 100,000 blocks, but switched to a static block reward of 10,000 DOGE after block 600,000. Ethereum took a different approach, with no hard limit and no set halving algorithm, but a design that planned for regular network updates. The block reward for Ethereum was set to 3 ETH in late 2017, and changed again to 2 ETH in early 2019.

If you're wondering about the discrepancy in block rewards, it's important to note that each coin algorithm also designates a different target rate for new blocks. Bitcoin aims for an average of 10 minutes per block—it's a long time, but it was the first cryptocurrency and once it gained popularity, it was increasingly difficult to change the underlying algorithm. For better or worse, Bitcoin has stuck with that 10-minute block time. Litecoin has a lot of similarities to Bitcoin, except it defined a block target time of 2.5 minutes, with the reward cut in half every 840,000 blocks.

Combined with the block reward halving structure, there's a limit to the total number of Bitcoins that will ever be created: 21 million, with half of them generated in the first 210,000 blocks, half of a half in the next 210,000, and so on. Technically, we'll never quite get there, but we'll get increasingly closer to the limit. Litecoin quadrupled the total number of coins it would produce to 84 million. Other coins have sometimes wildly differing schemes, but generally, scarcity is involved—rare things are worth more.

The critical factor with all coins is that the algorithms are designed to adjust the difficulty of finding a solution over time. The more hardware running the algorithms, the faster blocks would be found if the difficulty were static. The solution space ends up being adjusted to fit within certain parameters, which in turn resets the typical block time. This means that, for a given set of hardware, the more people running the mining algorithm, the fewer coins each one will get through mining over time. This has the side effect of creating an arms race to get more hardware up and running as quickly as possible, leading to ever-increasing mining farm sizes. That in turn leads to GPU shortages—or ASIC shortages (see "What are ASICs?" above).

Rewards aren't limited to just the block reward, though. Also included are all the fees for transactions included in the block. Early on, these fees tend to be very limited, but lately, as an example, the transaction fee rewards on Ethereum blocks are pretty close to matching the 2 ETH block reward, meaning each block is actually worth 4 ETH. The idea is that by the time the static block rewards approach zero, they've been superseded by the transaction fee rewards.

THE NEED FOR SPEED

Ultimately, coin mining ends up being a question of speed and efficiency. This generally gets measured in the hash rate, expressed using SI prefixes and H/s (hashes per second). The complexity of the hashing algorithms varies greatly, which means there's no universal hash rate for any set of hardware. Some algorithms are also designed to require more memory (such as

© WIKIMEDIA/ANTOINEBERCOVICI

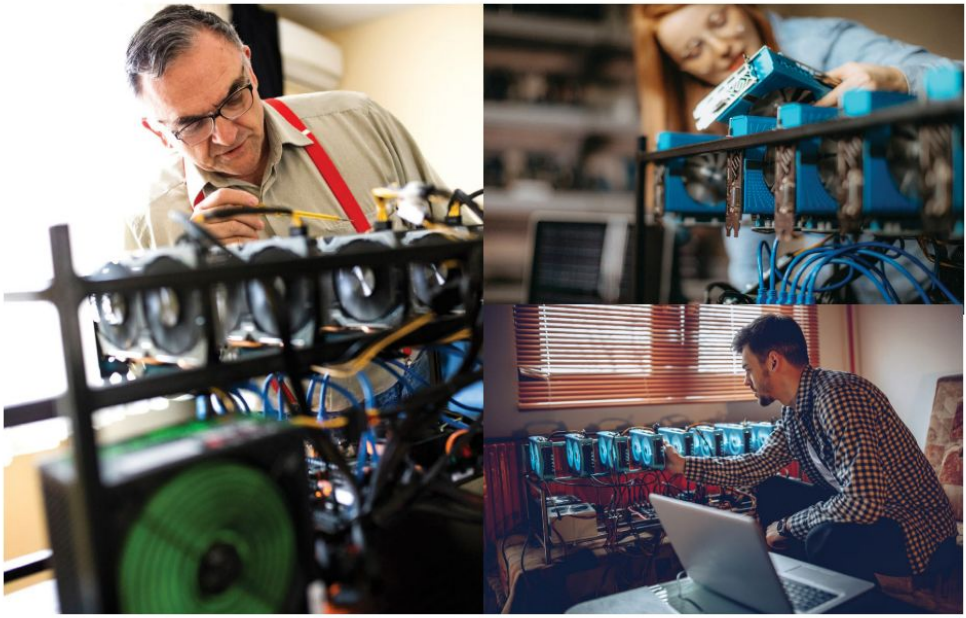

These examples are probably a bit more glamorous than your average mining setup, but this is what it could look like.

Ethereum), which can further limit the types of hardware that can effectively run the hashing calculations.

The fastest GPU right now, RTX 3090, has a hash rate of around 120MH/s (megahashes per second) for the Ethereum algorithm, called Ethash (an upgraded version of Dagger-Hashimoto, but we digress). The same GPU can do around 1,0000MH/s, or 10GH/s, with Bitcoin's SHA256 algorithm. Modern SHA256 ASICs, meanwhile, can do roughly 25TH/s.

Speed is important, but the amount of power used also factors into the picture. If some piece of hardware is 10 times faster but requires 10 times as much power, it's not really a net win. Tuning your hardware for optimal performance is another important task. Out of the box, for example, an RTX 3080 might use its full 320W of power to get around 85MH/s for Ethereum. Drop the GPU clock speed and limit the power while overclocking the memory and it's possible to get 95MH/s while only using 220W of power.

Lowering the power has the added benefit of lowering GPU temperatures, which means running fans slower and hopefully getting your GPU to last a lot longer. However, Nvidia's RTX 3080 and 3090 come with GDDR6X memory that can be more of a problem when it comes to thermals than the GPU cores. The reference 3080 and 3090 Founders Edition cards both hit 110 C on their GDDR6X memory at stock settings—that's not something we'd feel comfortable running as a long-term mining solution. Replacing the thermal pads, adding fans, and other mods can improve the situation, and we advise anyone doing mining with a GDDR6X card to use HWinfo64 to check memory temperatures.

HOW TO MINE CRYPTOCURRENCIES

The basics of mining cryptocurrencies haven't changed too much over the years. You need the appropriate hardware—GPU, ASIC, or even CPU—and the requisite software. There are a few different software packages available, many of which can be quite complex to configure and get running optimally. If you're looking for an easier path to dipping your toe into the mining waters, consider using NiceHash (see box over the page).

One thing we don't recommend is trying to mine "solo." That may have been possible in the early days of Bitcoin, or you could give it a shot if you build your own large-scale mining farm (please don't do that!), but these days you'll have far more success using a mining pool. That means picking a pool in the first place, and you'll want to research that via the Internet (for example, "best ethereum pool" should give you some reasonable options). Then you'll need to set up an account, download the appropriate mining software, and configure the launch settings as appropriate.

You'll want to pay close attention to any pool mining fees. These typically range from 0–3 percent, sometimes more, but plenty of good options usually land at the 1 percent value. It costs money to run the servers and infrastructure for a mining pool, and the fees go to support that (and make the pool owners a profit). Any "free" pools may end up being either unreliable, or with lower payouts that negate the benefit of not paying a fee.

One really important warning: Do not use the same password on any sites related to cryptocurrency mining. Seriously, just don't—and don't ask how we know this, please. Using a unique password on every site isn't just best security practice, it's a way to keep yourself from being robbed by potentially unscrupulous site owners.

WARNING: HIGH POWER

We've said a lot about the hows and whats of cryptocurrencies, but there's one final item to address, and it's a big one. While there are arguably worse ways to earn money than cryptocurrency mining, there are also a lot of better ways—ways that don't carry

© GETTY IMAGES/ VIKTORCVETKOVIC/ STEFANIKOLIC/ SOUTH AGENCY

MINING BTC WITH A GPU: NICEHASH

Arguably the easiest way to get started at mining cryptocurrencies is with NiceHash. We've noted that it's no longer profitable to mine Bitcoin directly using a GPU, because ASICs are so much faster and more efficient. That doesn't mean you can't get BTC from a GPU, however. NiceHash launched in 2014 and greatly lowered the barrier to entry for mining. It gets rid of the worry about what coin(s) to mine and is extremely easy to get running. You lease your PC's hashing power to other users, who get to choose what to mine, and you get paid in Bitcoin. NiceHash takes a small cut of the potential profits, and your PC can be up and mining in minutes.

The company's custom NiceHashMiner software benchmarks your hardware using various common hashing algorithms, and offers plenty of configuration options. NiceHashMiner then chooses whatever hashing algorithm currently pays the best rates, based on what people are willing to pay to lease your hardware. Sometimes a new coin launches, or someone wants to dedicate a lot of mining power at a specific coin, and they pay more to do so. Instead of mining Ethereum 24/7, your hardware will occasionally run some other algorithm, and it's all managed by the software. The software also has auto-update functionality to keep you on top of the latest developments.

There are other ways to mine with NiceHash, such as its QuickMiner software, which aims to make the process even easier. Alternatively, there's a custom NiceHash OS Linux-based install you can run in place of Windows. Performance and efficiency are potentially higher, and for dedicated 24/7 mining, it's supposed to be the best option.

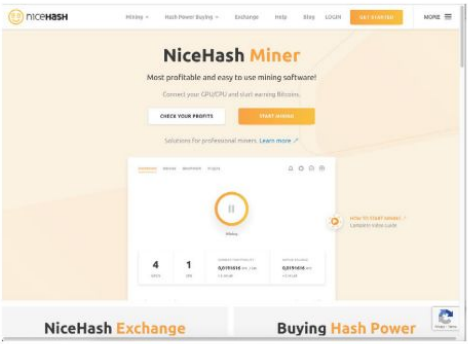

So, next time you hear someone say, "You can't mine Bitcoin with GPUs anymore," you can set them straight. While it's true that directly mining BTC with a GPU isn't a good idea, services such as NiceHash provide a convenient workaround. Alternately, just mine Ethereum or whatever on your own, and then trade that for Bitcoin.

The Bitcoin network uses 3,000 times as much energy as the world's fastest supercomputer, the Japanese Fugaku.

© STR / JIJI PRESS / AFPJ / JAPAN OUT (PHOTO BY STR/JIJI PRESS/AFP VIA GETTY IMAGES, NICEHASH

nearly the volatility and risk of coin mining. Procuring all of the necessary mining hardware can take time and a lot of money, which makes it difficult for PC enthusiasts to upgrade their PCs. Hardware can wear out sooner than you expect. But the bigger issue, by far, is that mining puts a ton of computing power to the task of securing the blockchain.

Best case, using the most efficient hardware, the Ethereum network would currently use about 1 billion watts of power, and Bitcoin would use 5.5 billion watts. Unfortunately, it's actually a lot more in both cases, because a large chunk of the hashing isn't done using the most efficient hardware. Online estimates peg the power use of the Ethereum network at around 21TWh per year, and 45kWh per transaction. Put another way, Ethereum hashing uses 60GWh every day, which costs around $6 million. That works out to $4.50 in power costs (using $0.10 per kWh) just to send Ethereum from one wallet to another.

Bitcoin is even worse. The network currently uses an estimated 78TWh per year, with a cost for a single transaction of around 690kWh. That's about as much power as a typical US household uses in 24 days! That means Bitcoin uses nearly $8 billion per year just keeping the blockchain secure. That figure continues to grow, and while Bitcoin and Ethereum are the two largest cryptocurrencies, there are dozens of others that together require another large chunk of power.

Looking at the costs and power going into these networks, it's difficult to remain optimistic about the long-term potential. We're strip-mining digital coins, and that's unsustainable. That's part of the cyclic nature of Bitcoin, where it finally reaches a level that can't be sustained, companies and people stop investing in it, and there's a massive correction. A few years later, enough people have forgotten the last time this all happened that we see a repeat. There are other proposed alternatives to current reward schemes, such as proof-of-stake instead of proof-of-work, but they have their own drawbacks.

Imagine if all of the effort being put into cryptocurrency mining was instead put into research. While Bitcoin's SHA256 hashing isn't useful for many other tasks, the most powerful supercomputer in the world—Japan's Fugaku—has a peak power consumption of around 30MW. That's a modest city's worth of power, and yet the Bitcoin network uses roughly 3,000 times as much energy. What could scientists do with another 3,000 of the fastest supercomputers? We don't know, but we'd love to find out. The money that's pouring into cryptocurrencies could be put to far better uses.

These drawbacks won't stop people from mining, though, at least not without help. Governments have started to crack down on mining farms, however—China's Inner Mongolia recently issued a ban on cryptocurrency mining. Other areas of the world may follow suit. Regulation seems inevitable, and for something that's supposed to free people from the oppression of the banking system, it's difficult to get past the fact that every coin eventually derives its worth from how much it is worth in dollars, pounds, yen, and so on.

If you still want to give mining a shot, particularly if you already have a modest PC and you live in an area where power isn't too expensive, or the power grid isn't already having severe issues, we won't object too much. But please don't bet the farm or your retirement on PC hardware and upgrades solely used in pursuit of mining. There's a real likelihood of chasing diminishing profits, potentially for years. You have to keep on mining in the hope of striking it rich with a future surge in cryptocurrency prices. We doubt we've heard the last of Bitcoin or Ethereum mining, but there's a good chance anyone trying to start now has already missed the current wave.

All that power has to come from somewhere, and electricity generation isn't always the greenest of industries.

© GETTY IMAGES/PAUL SOUDERS

FEED YOUR MIND

www.howitworksdaily.com

HOW IT WORKS™

FEATURING NEW AUGMENTED REALITY EXPERIENCE!

BUY YOUR ISSUE TODAY

Print and digital subscriptions available at

www.magazinesdirect.com

How It Works magazine @HowItWorksmag howitworks@futurenet.com howitworksmag

SUBSCRIBE TO MAXIMUM PC

THE GO-TO RESOURCE FOR PC ENTHUSIASTS...

MAXIMUM PC
VIRTUAL REALITY
The state of play in 2021
Tested: The best headsets you can buy
Oculus versus Vive: Which is the true virtual king?
A deep dive into VR and game development
CRYPTO-MINING BOOM OR BUST!

1

THE DIGITAL EDITION
ONLY $18.00
EVERY YEAR

Instant digital access on your iPad, iPhone, and Android device.

MAXIMUM PC
VIRTUAL REALITY
The state of play in 2021
Tested: The best headsets you can buy
Oculus versus Vive: Which is the true virtual king?
A deep dive into VR and game development
CRYPTO-MINING BOOM OR BUST!
STEP-BY-STEP GUIDE

2

THE PRINT EDITION
ONLY $30.00
EVERY YEAR

Every issue delivered to your door for a fraction of the store price.

SUBSCRIBE ONLINE TODAY

http://bit.ly/MaxpcD21

A POTTED HISTORY OF LINUX

Or how a 21-year-old's bedroom coding project took over the world and a few other things along the way. *By Neil Mohr*

LINUX ONLY EXISTS because of Christmas: On January 5, 1991, a 21-year-old computer science student, who was currently living with his mom, trudged through the (we assume) snow-covered streets of Helsinki, with his pockets stuffed full of Christmas gift money. Linus Torvalds wandered up to his local PC store and purchased his first PC, an Intel 386 DX33, with 4MB of memory and a 40MB hard drive. On this stalwart machine he would write the first ever version of Linux. From this moment on, the history of Linux becomes a love story about open-source development, software freedom, and open platforms.

Previous to walking into that computer store, Linus Torvalds had tinkered on the obscure UK-designed Sinclair QL (Quantum Leap) and the far better known Commodore Vic-20. Fine home computers, but neither was going to birth a world-straddling kernel. A boy needs standards to make something that will be adopted worldwide, and an IBM compatible PC is a good place to start. But we're sure Torvalds' mind was focused more on having fun with *Prince of Persia* at that point than developing a Microsoft-conquering kernel.

Let's be clear: A 21-year-old, barely able to afford an Intel 386 DX33, was about to start a development process that would support a software ecosystem, which in turn would run most of the smart devices in the world, a majority of the Internet, all of the world's fastest supercomputers, chunks of Hollywood's special effects industry, SpaceX rockets, NASA Mars probes, self-driving cars, and whole bunch of other stuff besides. How the heck did that happen?

© MAGICTORCH

TO UNDERSTAND HOW Linux got started, you need to understand Unix. Before Linux, Unix was a well-established operating system standard through the 1960s into the 1970s. It was already powering mainframes built by the likes of IBM, HP, and AT&T. We're not talking small fry, then—they were mega corporations selling around the globe.

If we look at the development of Unix, you'll see certain parallels with Linux: Freethinking academic types who were given free rein to develop what they want. But whereas Unix was ultimately boxed into closed-source corporatism, tied to a fixed and dwindling development team, eroded by profit margins and lawyers' fees, groups that followed Linux embraced an open approach, which enabled free experimentation, development, and collaboration on a worldwide scale. Yeah, yeah, you get the point.

Back to Unix, which is an operating system standard that started development in academia at the end of the 1960s as part of MIT, Bell Labs, and AT&T. The initially single or uni-processing OS, spawned from the Multics OS, was dubbed Unics, with an assembler, editor, and the B programming language. At some point, that "c" was swapped to an "x," probably because it was cooler, dude.

Later, someone needed a text editor to run on a DEC PDP-11 machine. So, the Unix team obliged and developed roff and troff, the first digital typesetting system. Such unfettered functionality demanded documentation, so the "man" system (still used to this day) was created with the first Unix Programming Manual in November 1971. This was all a stroke of luck, as the DEC PDP-11 was the most popular mini-mainframe of its day, and everyone focused on the neatly documented and openly shared Unix system.

Ken Thomas (left) and Dennis Ritchie are credited with creating much of the original UNIX, while Ritchie also developed the C language as well.

In 1973, version 4 of Unix was rewritten in portable C, though it would be five more years until anyone tried running Unix on anything but a PDP-11. At this point, a copy of the Unix source code cost almost $100,000 in current money to license from AT&T, so commercial use was limited during the 1970s. However, moving into the 1980s, costs rapidly dropped, and widespread use at Bell Labs, AT&T, and among computer science students propelled the use of Unix. It was considered a universal OS standard, and in the mid-1980s, the POSIX standard was proposed by the IEEE, backed by the US government. This makes any operating system following POSIX at

Linus Torvalds, being interviewed by sister magazine *Linux Format* in 2012.

© WIKIMEDIA

Ian Murdock, founder of Debian, being interviewed by *Linux Format* in 2007.

least partly if not largely compatible with other versions.

At the end of the 1980s, the Unix story got messy, with lots of commercial infighting, competing standards, and closing off of standards, often dubbed Unix Wars. So, while AT&T, Sun Microsystems, Oracle, SCO, and others argued, a certain Finnish boy was about to start university.

WE GNU THAT

Before we dive into the early world of Linux, there's another part of the puzzle of its success that we need to put in place: the GNU Project, established by Richard Stallman. Stallman was a product of the 1970s development environment, a freethinking, academic, hippy type. One day, he couldn't use a printer, and because the company refused to supply the source code, he couldn't fix the issue—supplying source code was quite normal at the time. He went apoplectic and established a free software development revolution: an entire free OS ecosystem, free software license, and philosophy that's still going strong. Take that, proprietary software!

The GNU Project was established by Stallman in 1983, with GNU being a hilarious (to hackers) recursive acronym for "GNU is Not Unix." Geddit? Its aim was to establish a free OS ecosystem with all the tools and services a fully functioning OS requires. Bear in mind, most of the tools created then are still used today. By 1987, GNU had established its own compiler, GCC, the Emacs editor, the basis of the GNU Core Utilities (basic file manipulation tools like list, copy, delete, and so on), and a rudimentary kernel. But just as importantly, Stallman had cemented his ideal of software freedom with the "copyleft" GPL software license, and his manifesto setting out the four software freedoms enabling users to run, study, modify, and distribute any software, including the source, for any reason.

The GPL remains the strongest copyleft license, and while it has perhaps fallen out of vogue, it's still regarded as the best license for true open-source development, and cements most Linux distros. GCC is still an industry standard, Emacs remains a feature-rich development environment, and the GNU Core Utilities are still widely used in certain POSIX systems and most Linux distros.

You could argue that without the GNU Project being established, Linux would never have taken off. The GPL license (adopted early on in Linux development) forces all developers to share back their enhancements to the source code. It's a feedback loop that promotes shared improvements. Alternative open-source licenses enable corporations to take source code and never share back improvements, meaning the base code is more likely to remain static. This was backed by a generation of developers that grew up studying and using Unix, looking for a truly freed open-source OS to contribute to.

KERNEL FREAX

We're getting ahead of ourselves. Linus Torvalds had his Intel 386, was studying computer science at the University of Helsinki, and was using the MINIX 16-bit OS and kernel. MINIX is a POSIX-compatible Unix-like OS and micro-kernel. In 1991, it had a liberal license, costing just $69, offering the source code but restricted modification and redistribution.

We imagine the 16-bit limitation spurred Torvalds to create his own 32-

BIG BUSINESS VS. LINUX

Being the root of evil and all, whenever money is involved, things can turn nasty. So, when the big players in the enterprise and business markets began to see Linux distros as a threat, lawyers were called.

A series of leaked Microsoft memos from August 1998, known as the Halloween Documents for the date they were released, detailed Microsoft's private worries that Linux, and open-source development in general, was a direct threat to its business, along with ways to combat its uptake. This private view was in direct conflict with the company's public line on the matter; though Steve Ballmer famously called Linux a cancer in 2001. The documents are available at www.catb.org/~esr/halloween, and in them Microsoft predicted that "Linux is on track to eventually own the x86 UNIX market…." It was correct.

There was little Microsoft could do to combat Linux, as it couldn't be bought. The documents suggested extending open protocols with Microsoft's own proprietary extensions (that didn't work), and seeding the market with fear, uncertainty, and doubt (FUD) also failed.

There was another angle, however: Help a company that's suing over copyright infringement of the source code. In 2003, a company called SCO claimed part of its UNIX System V source code was being used within Linux, making it an unauthorized derivative of UNIX. SCO sued IBM for $1 billion (among many other companies), and demanded end users pay a Linux license fee. Microsoft leaped into action and paid SCO $106 million, as detailed in a leaked and verified SCO memo. After years of legal arguments, a code audit found there to be no evidence of copied UNIX code in the Linux kernel. SCO went bankrupt in 2009.

© GETTY IMAGES

Richard Stallman, the creator of the GNU free software movement, being interviewed by *Linux Format* in 2011.

bit kernel, but he states the license restrictions were also key. So, on August 25, 1991, Linux posted to comp.os.minix that he was developing his own free OS, "nothing professional like GNU," and it'd only support AT disks, as that's all he had.

This was developed on a MINIX system, compiled on GNU GCC, and he'd ported GNU bash. Torvalds had planned to call his OS Freax, combining "Free," "Freak," and "Unix," but once he'd uploaded it to ftp.funet.fi, a volunteer admin (Ari Lemmke) renamed it Linux, as he thought it was better. So, version 0.01 of Linux was released to the world in September 1991. One telling part of the release notes states: "A kernel by itself gets you nowhere. To get a working system you need a shell, compilers, a library, etc... Most of the tools used with Linux are GNU software and are under the GNU copyleft. These tools aren't in the distribution—ask me (or GNU) for more info."

Importantly, this outlines Linux's reliance on other GPL-licensed tools, and shows the use of the term "distribution," now shortened to "distro." As Torvalds points out, an operating system isn't a kernel alone; it's a collection of tools, scripts, configs, drivers, services, and a kernel, lumped together in an easier form for users to install and use.

As for the license, Torvalds initially used his own, which restricted commercial use, but by January 1992, he'd been asked to adopt the GPL, and had stated the kernel license would change to align it with the other tools being used. It was December 1992, and for the release of v0.99, the Linux kernel was GPLv2-licensed. This cemented the legal clause that anyone using the kernel source has to contribute back any changes used in published code.

We're not going to follow the development of the Linux kernel itself

Thankfully, by being buggy, SoftLanding Linux kick-started some good distros.

The original Raspberry Pi model always ran a version of Linux.

© KEVIN NIXON, LINUXCENTER.RU

in detail, but how Linux distros from this point were developed and branched into a wide-ranging ecosystem. This enabled the use of the Linux kernel to grow rapidly, as distros not only ease the installation of Linux (which early on was a complex process of source compilation, gathering the right tools, creating filesystem layouts by hand, and bootloaders, all from the Terminal on systems with limited resources), but one distro can also become the base for a whole new distro, tailored for a new use or audience.

PRIMORDIAL SOUP

As v0.01 was only released in September 1991, the first distribution of Linux—though by modern standards, it's lacking in every department—created by H. J. Lu, was simply called Linux 0.12. Released at the end of 1991, it came on two 5.25-inch floppy disks, and required a HEX editor to get running. One disk was a kernel boot disk, the other stored the root OS tools.

In those early days of distro evolution, things changed rapidly. Development was quickly adding functionality, and people tried out the best ways to package a Linux-based OS. MCC Interim Linux was released in February 1992 with an improved text-based installer, and was made available through an FTP server. X Windows—the standard Unix windowing system—was ported, and TAMU Linux was released in May 1992 with it packaged: The first graphical distro.

Whilst all of these are notable as being among the first Linux distros, they didn't last. The same can be said for Softlanding Linux System (SLS), also released in May 1992, which packaged X Windows and a TCP/IP network stack. It's notable, though, because its shortcomings (bugs and a change to the executable system) inspired the creation of the two longest-running and, in many ways, most influential Linux distros: Slackware and Debian.

EVOLUTION OF A GENUS

Now a number of base distros appear, reliably maintained by individuals, groups, or businesses. Once they're established, stable, and become popular, offshoots branch from these root distros offering new specializations or features. This creates a number of base distro genera, formed around the original package manager and software repositories.

The effect is a Linux family tree, where you can date all distros back to an initial root release. Some branches sprout and die off; either the group maintaining it disbands or there's no wider interest. Some branches become so popular they create a whole new genus, becoming the basis for a further expansion of distros.

As with plants and animals, offshoots inherit traits, the base install, package manager, and software repositories being key. A package manager is how the OS installs, updates, removes, and maintains the installed software, which includes downloading software packages from the managed software servers, called repositories. This can get contentious, as these child distros are leeching off the parent's bandwidth, but initially, while they're growing, this use won't look much different from normal user activity.

Bear in mind we're back in 1992—you're lucky if people have 14.4Kb/s dial-up modems at home, while expensive T1 lines (1.54Mb/s) are limited to academic institutions and larger businesses. The early TAMU v1.0 distro required 18 disks for the 26MB binaries, and 35 disks for the 50MB compressed (200MB uncompressed) source code. This obviously limited access in these early days to academics and those in suitable businesses, so distro evolution was slow.

MEET THE ANCESTORS

Softlanding Linux System was popular, but it was buggy and badly maintained,

All the world's fastest supercomputers are now running Linux.

© OAK RIDGE NATIONAL LABORATORY

THAT'S NOT LINUX!

There have been plenty of other kernels and operating systems, many open source, but all languishing in the footnotes of history. The original UNIX releases had their day, but the POSIX standard they established, the shackles and legalities of overbearing license restrictions, and the desire of coders to work on freed code were all too much.

Minix could have been its own Linux, being the inspiration for it, and it is still maintained, but again original license limitations and possibly a lack of ambition meant it was always more of an academic curio.

We haven't yet mentioned the **Berkeley Software Distribution** (BSD), which is a success today in its own right—it powers the modern Apple Darwin kernel—and followed Linux in many steps, shares a lot of the same libraries and tools, but was tied up in legalities in the early 1990s, which probably allowed Linux to gain the weight of developer focus.

BeOS is the best OS you've (probably) never heard of. Launched in 1995, it absolutely shamed Windows 95 and macOS at the time. Initially developed on PowerPC systems, the idea being to replace macOS, it never took off and ceased development around 2001. Of course, www.haiku-os.org is an open source implementation of the BeOS API, so it lives on.

MikeOS (http://mikeos.sourceforge.net) is an x86 OS written in assembly by ex-*Linux Format* disc editor Mike Saunders. It's an open-source, 16-bit, real-mode OS aimed at teaching how an OS works.

Last but not least, **Temple OS** (http://templeos.org) is an incredible public domain work by coder Terry A Davis, sadly deceased. Suffering from schizophrenia, Davis was compelled to write HolyC and Temple OS.

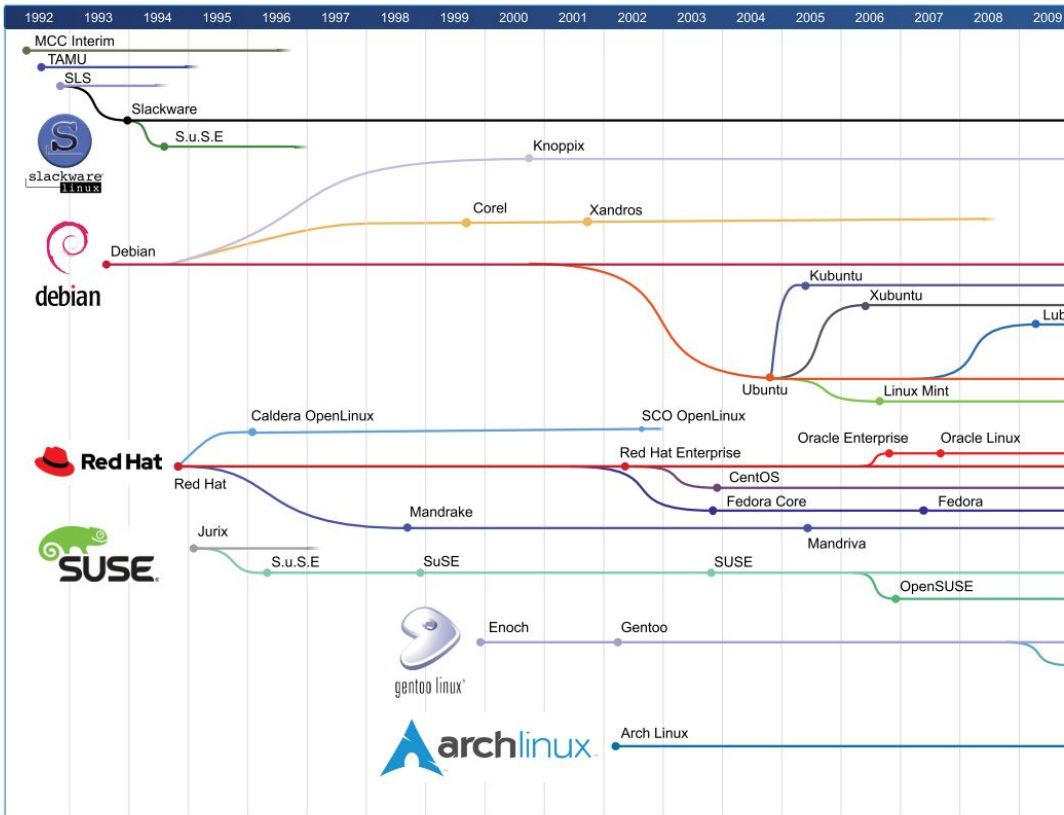

From its birth in 1991, Linux has been a force to be reckoned with. Its many offshoots are still going strong as it hits its thirtieth anniversary.

so in July 1993, Patrick Volkerding forked SLS and created Slackware—so named because it wasn't a serious undertaking at the time, and was a reference to the Church of the SubGenius—this is the oldest Linux distro still maintained, and it's about to see its version 15 release after 28 years. Slackware is interesting because it's very much controlled and maintained by Volkerding, while followed by a small but enthusiastic band of users and contributors. Whereas many other distros have taken on modern enhancements, Volkerding sticks to older more traditional "Unix" ways of controlling services on Slackware. There's no formal bug tracking, no official way to contribute to the project, and no public code repository. This all makes Slackware very much an oddity that stands on its own in the Linux world. Due to its longevity, however, Slackware has attracted a couple dozen offshoots, and at least half are still maintained today.

In August 1993, Ian Murdock, also frustrated by Softlanding Linux System, established Debian, a combination of "Debby," his girlfriend's name at the time, and "Ian." From the outset, it was established as a formal, collaborative open project in the spirit of Linux and GNU.

Early on in the Debian project, Bruce Perens maintained the base system. He went on to draft a social contract for the project and created Software in the Public Interest, a legal umbrella group to enable Debian to accept contributions. At the time, Perens was working at Pixar, so all Debian development builds are named after *Toy Story* characters. The Debian logo also has a strong similarity to the mark on Buzz Lightyear's chin.

Debian is arguably the single most influential and important Linux distro ever. Just the sheer number of branches of distros from it would attest to that, but Debian is renowned for its stability, high level of testing, dedication to software freedom, and being a rigorously well-run organization. It's testament to its creator, Ian Murdock, who sadly passed away in December 2015.

GETTING COMMERCIAL

Things were still moving slowly into 1994—just a single Slackware fork called SUSE and a few random Linux sprouts appeared, but all died out. In October 1994, Red Hat Linux was publicly released. Red Hat was established as a for-profit Linux business, initially selling the Red Hat Linux distribution and going on to provide support services. Red Hat went public in 1999, achieving the eighth biggest first-day gain in the history of Wall Street. It entered the NASDAQ-100

© WIKIMEDIA/MANTICORE010/EDITED VERSION

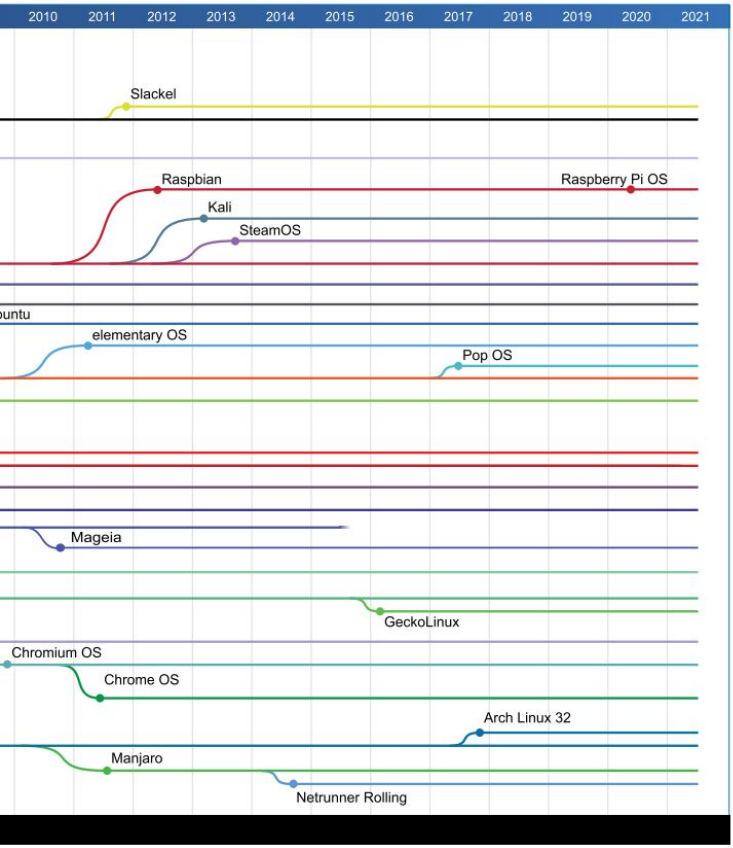

in December 2005 and topped $1 billion annual revenue in 2012. IBM purchased Red Hat in October 2018—24 years after its first release—for $34 billion. So that worked out very well.

Red Hat Linux was relaunched as Red Hat Enterprise in 2001, and its commercial success attracted a wide range of forks. Notably, Red Hat directly supports Fedora as its testing distro and CentOS as its free community edition. Or it did. CentOS is being shuttered—to understandable community disdain—and a rolling release CentOS Stream is replacing it. As an alternative, Red Hat Enterprise is now offered freely to community projects with fewer than 16 servers.

Meanwhile in Germany, SUSE (Software und System Entwicklung) started life as a commercially sold German translation of Slackware in late 1992. In 1996, an entire new SUSE distro and business was released, based on the Dutch Jurix Linux, selling the new distro and support services. SUSE was purchased by Novell in 2003, and in 2005, the openSUSE community edition was launched, while SUSE Linux Enterprise was developed in tandem for its commercial arm. SUSE was acquired in 2018 for $2.5 billion and returned double-digit growth through 2020, with a revenue of over $450 million. Despite its success, SUSE and openSUSE have only ever attracted a couple of forks. Possibly down to their European roots.

DISTRO INFERNO

Between the creation of Red Hat in 1994 and 2000, there were a good number of Red Hat spin-offs, because at that point there was clear commercial interest in Linux. Through this period, Linux was best suited to business server tasks, where much of the open-source Unix work had been focused. However, by the end of the 1990s, 56k modems had become commonplace, early home broadband was just appearing, and modern graphical desktops were in development. Linux was about to get a whole new audience.

One early example was Mandrake Linux, in mid-1998. A fork of Red Hat, it was crazily aimed at making Linux easy to use for new users, using the new Kool Desktop Environment (KDE). The French/Brazilian development team gained a lot of attention but, ultimately, financial problems closed the project in 2011. However, its spirit continues in the excellent but less well known Mageia and OpenMandriva projects.

With Mandrake pointing the way, the early 2000s saw an explosion of distro releases. With the Debian project at this point being well established, well regarded, and well known, it became the basis for hundreds of Linux distros. But we'll only mention one: Ubuntu, released in 2004 by South African millionaire Mark Shuttleworth, who jokingly calls himself the self-appointed benevolent dictator for life. The Ubuntu Foundation was created in 2005 as a philanthropic project—Ubuntu is a Zulu word meaning humanity—to provide quality open-source software, with Canonical as the supporting commercial arm.

Ubuntu as a branch of Debian has itself seen over 80 distros branched from it, while Ubuntu has the highest share of all desktop Linux installs—though this is notoriously hard to measure—when users are polled. Why Ubuntu became so popular is hard to pinpoint. Like Mandrake before it, it set out to make desktop Linux easy for first-time users. It also offered the distro on free CDs via its ShipIt service until 2011, alongside fast, reliable server downloads. Plus, it was based on the popular Debian, it jumped on the new, slick Gnome desktop, and it set out a regular six-month release cycle, with a Long Term Support release every two years. Support was for 18 months (now nine months) for regular releases, and 36 months for LTS ones (now five years).

It also offered great forums and help sites, along with a community council, and support for forks such as Xubuntu, Lubuntu, and many others. It had sane defaults, too, and made it easier to install display drivers (an absolute pain 10-plus years ago), while offering a huge catalog of tested, ready-to-run open-source software and dedicated server builds. We guess when you say it out loud, it sounds pretty compelling.

Two core release branches we'll quickly mention are Arch Linux and Gentoo,

techradar.

The home of technology

techradar.com

both released around 2000. Gentoo (named after the fastest penguin in the world) is a built-from-source distro compiled with specific optimizations for the hardware it's going to run on. This is very clever but also very time-consuming. Google Chrome OS is derived from Gentoo. In early 2002, Arch Linux was released, devised as a minimalist distro, where the user does much of the installation work to create an OS with just the parts required. This DIY approach was partly why Arch is renowned for its amazing documentation and for rolling out the earliest release of new versions of software.

MODERN CONSOLIDATION

At the height of the distro madness (around 2010), there were almost 300 Linux distros, an unsustainable number, with many just repeating basic desktop functionality already available in core root distros. Progressing into the 2000s, and with increasing complexity in maintaining a modern OS, the number of Linux distros started to reduce, but that didn't stop well-organized groups creating popular new distro forks when they felt a need.

A good example is Raspberry Pi OS, a rebrand of Raspbian, itself a fork of Debian. The new Arm-based hardware platform needed a dedicated operating system, so picking up Debian and refitting it for the Raspberry Pi, including educational software, libraries for its GPIO access, and tailored tools to configure its hardware, made absolute sense.

Linux hardware specialist System76 was tired of niggling software issues associated with using other distros, and wanted direct control. So, it introduced Pop!_OS, a fork of Ubuntu, to not only directly support its laptops and desktop hardware, but also its customers' needs. It's a slick, modern distro, with support for popular software and hardware.

Linux Mint started in 2006 as a small personal Ubuntu fork project. When Ubuntu changed to its "modern" Unity desktop design in 2011, many users revolted. The Linux Mint project created its own "classic" desktop, called Cinnamon, in 2012, and it brought many former Ubuntu users with it. The Linux Mint project has stuck with its "user first" design approach, and evolved well.

This doesn't even touch upon commercially focused distros, such as Android, Chrome OS, Intel's ClearOS, Google's Wear OS, Sailfish OS, and the host of server-specific distros. Even today, there are well over 200 (we counted 265) active Linux distros, and they're as diverse, interesting, and wonderful as the communities that use them.

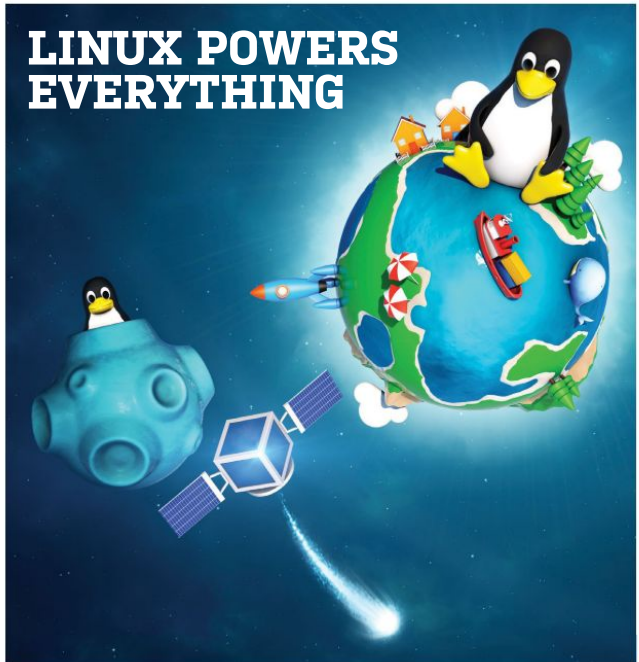

Developing software for supercomputers is expensive; during the 1980s, Cray was spending as much on software development as it was on its hardware. In a trend that would only grow, Cray initially shifted to UNIX System V, then a BSD-based OS, and eventually, in 2004, SUSE Linux to power its supercomputers. This was matched across the sector, and the top 500 supercomputers (www.top500.org) now all run Linux.

Internet services have also all been developed to run on Unix systems. Microsoft and BSD systems do retain a good slice of services, but over 50 percent of web servers are powered by Linux. Recent moves to virtual services with container-based deployment are all Linux-based. Microsoft's cloud service Azure reports Linux as its largest deployment OS and, more to the point, Google uses Linux to power most of its services. As do many other service suppliers.

Android's mobile OS share dropped in 2020 to just 84 percent—it's powered by Linux. Google bought the startup that was developing Android in 2005. LineageOS (https://lineageos.org) is a well-maintained fork of Android and supports most popular handsets well after their manufacturers abandon them.

Space was thought to be Linux's final frontier, because it's not a certified deterministic OS, which is the gold standard for real-time OSes in mission-critical situations. Turns out SpaceX rockets use Linux to power their flight systems, using a triple-redundancy system, while NASA has sent Linux to Mars in its latest drone. Tesla is also reportedly running Linux in its cars.

Linux has also been at the heart of Hollywood's special effects since 1997's *Titanic* used a Linux server farm of DEC Alphas at Digital Domain to create its CGI. DreamWorks' *Shrek* in 2001 was the first film entirely created on Linux systems. Meanwhile, Pixar ported its Renderman system to Linux from SGI and Sun servers around 2000, in time to produce *Finding Nemo* in 2003.

© MAGICTORCH

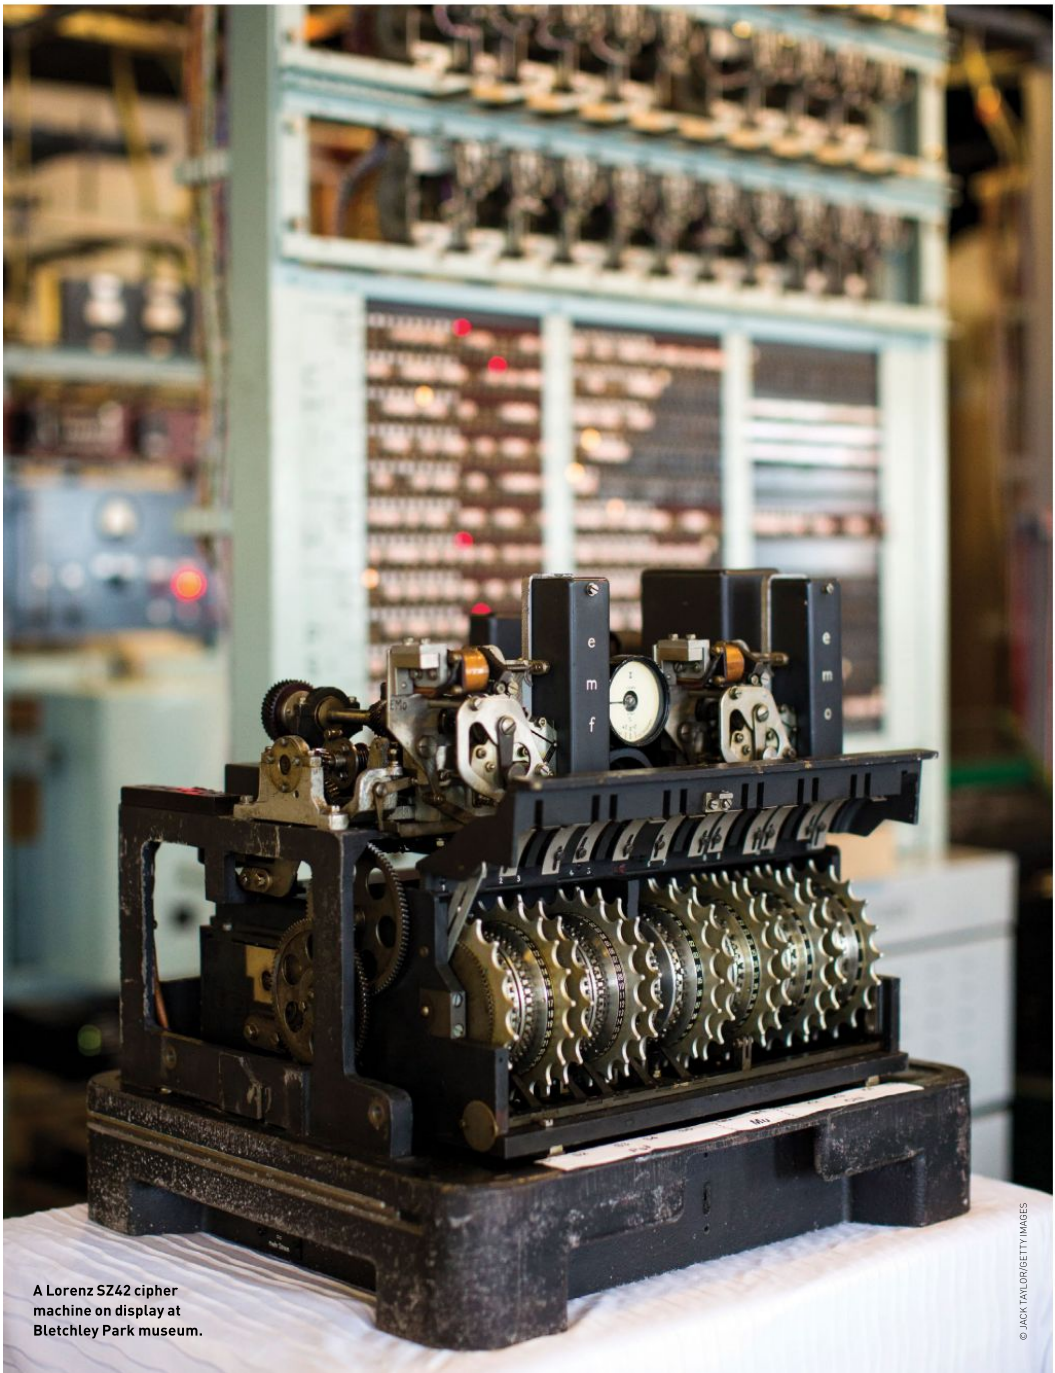

A Lorenz SZ42 cipher machine on display at Bletchley Park museum.

© JACK TAYLOR/GETTY IMAGES

A COLOSSAL SECRET

Ian Evenden discovers how British WWII braniacs, faced with an impossible task, helped save millions of lives and invented the computer as we know it

Your country is at war, and you start picking up directional radio signals that sound like nothing you've ever heard before. The enemy is highly mechanized, makes use of the latest technology, and has a record of using complex codes and ciphers in its transmissions. This is something new, however. Instead of the dots and dashes of the Morse code you're used to intercepting, this sounds like a harsh wailing.

These days, we'd probably equate the sound with a 56K modem, but back in the 1940s, it was known as teleprinter code. Nazi high command, not content with the "unbreakable" Enigma cipher it used to spread orders among its companies and brigades, was also using a machine for higher-level communications called the Lorenz SZ that British codebreakers had never seen, and would not see until the end of World War II. Despite this, they were able to deduce how it worked, and crack the encipherment.

This is the story of a country house in England full of mathematicians, crossword enthusiasts, the occasional genius, and a man from the postal service more used to creating automated telephone exchanges. Together, they read German military communications and shortened WWII by up to four years, saved possibly millions of lives, and created Colossus, the first programmable electronic computer, in the process.

This is a complicated story, and we've had to leave bits out, but bear with us, there's still a lot of backstory. The tale of how the British, Alan Turing in particular, cracked Enigma, opening up Nazi communications, and making it possible to anticipate their every move is already fairly well known, but it has some parallels with the decryption of the Lorenz machine that can't be ignored. The Nazis had mechanized the sending of secret messages, so a mechanized approach was needed for reading them. Lorenz had similar weaknesses to the Enigma, especially if you could be confident of the content of certain parts of messages, such as ending everything with "Heil Hitler," for example.

Enigma was an electro-mechanical machine with three, sometimes four, rotor wheels and a plugboard. Input came from a keyboard, and its output was a lamp that lit up behind a grid of letters. The rotors moved every time a key was pressed, one full rotation triggering the next one to move, then the next. The plugboard swapped letters over, so connecting A to S always switched those letters. With three rotors from a set of five in use, and 10 connections on the plugboard, the machine was capable of nearly 159 quintillion (18 zeroes) combinations. Adding a fourth rotor, as the German Navy did, just made things worse.

Lorenz had 12 rotors, and acted as an add-on to a teleprinter machine (you type at one end, the message comes out printed on paper tape at the other), but otherwise operated in a similar way. Once the machine's settings for the day were entered, however, it was much more automatic. The enciphered text was encoded as a 5-bit Baudot code (International Telegraph Alphabet No. 2) for transmission, which made use of a modulated on/off signal—binary, in other words—that was the source of the modem-like screeching the British listening stations were picking up.

"It was a binary machine, but not understood as a binary machine," says Paul Gannon, author of *Colossus: Bletchley Park's Greatest Secret*. "They used a different terminology at the time—mark and space. They didn't think of it as zero and one; that's a backward projection from what we know now."

The Lorenz machine then encrypted the Baudot further with a Verman stream cipher. Ten of the 12 rotors in the machine generated a key stream in two groups of five, which was combined with the plaintext using a logical operation known as Exclusive Or (XOR—see table on page 51). The remaining two wheels added a pseudorandom "stutter" to the key stream, complicating it further. Each rotor had a different number of cams on it, which could be raised or lowered to create the wheel settings. The total number of settings is a number so large it would be followed by 150 zeroes (there's 18 in a quintillion, remember).

"There's no reason why it had to be like this," says Gannon, "but with Enigma they would write the message down on a piece of paper, one guy would then read out the message, another would type in each letter of the message and read off the encoded letter, then write it down. Then somebody else would transmit that using Morse code. At the other end, they would read the Morse code, and reverse the message with an Enigma set up with the same settings. But the teleprinter was a central communications tool within Germany, and especially with the Nazis when they took over, and so they wanted to carry on using it. First of all they transmitted from teleprinter to teleprinter over cables, and then using a wireless link."

Gannon continues: "Your message was automatically encrypted and at the other end automatically decrypted, so all you needed was one typist at each end, rather than three people writing things down. Nobody saw the encrypted version of the message. The teleprinter was the key to communications within the German army, and this suited the British at the time because they had cracked the German naval and air force Enigma, but hadn't cracked the army Enigma. So when they started intercepting these weird sounds, which they described as being like a cow wailing, it took some time to work out what they were."

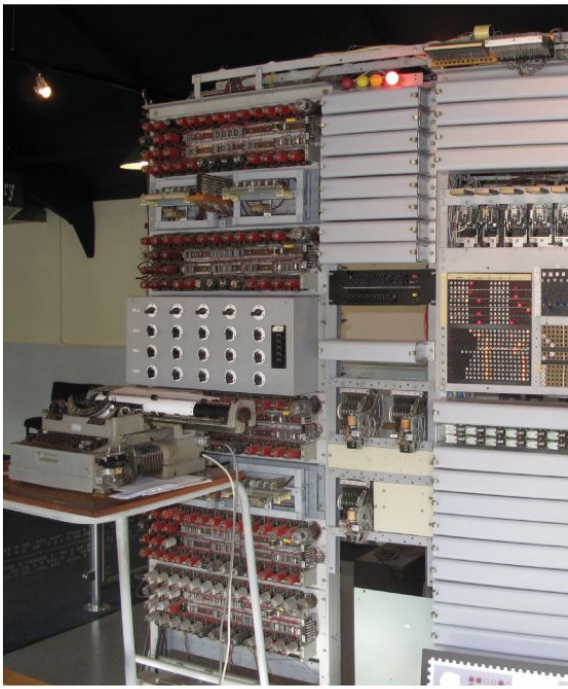

Frontal view of the reconstructed Colossus at The National Museum of Computing, Bletchley Park, UK.

The first transmissions were intercepted in 1941, a link between Berlin, Germany, and Athens, Greece. Used to dealing with Morse code, the British classified these as NoMo, or non-Morse, and the Berlin-Athens link picked up the code name Tunny (tuna fish). Tunny would become the code name for the Lorenz machine itself (*Geheimschreiber*, or secret writer, in German) after the link closed in 1942, while the teleprinter operation as a whole became Fish. As new links sprang up, they attracted new species of sea life: Berlin to Tunis was Herring, Berlin to Paris Jellyfish, and Berlin to Oslo was given the code name Mullet. This was a recurring theme—three-rotor Enigma transmissions were code-named Dolphin, while four-rotor transmissions were Shark. The intelligence gathered by codebreaking was known as Ultra, and was an incredibly closely guarded secret.

Bletchley Park, home of the innocuously named Government Code and Cipher School (think today's NSA), is a country house in the middle of England, built in the 1880s. These days, it's a museum, but during WWII, it was a branch of the British intelligence services. Its activities were so highly classified that it wasn't until 1974, with the publication of a book, *The Ultra Secret*, by F. W. Winterbotham, that details began to trickle out.

What took slightly longer to emerge, however, were the details of what they'd built there. The Enigma cipher was broken through human ingenuity (a lot of early work had been done in Poland

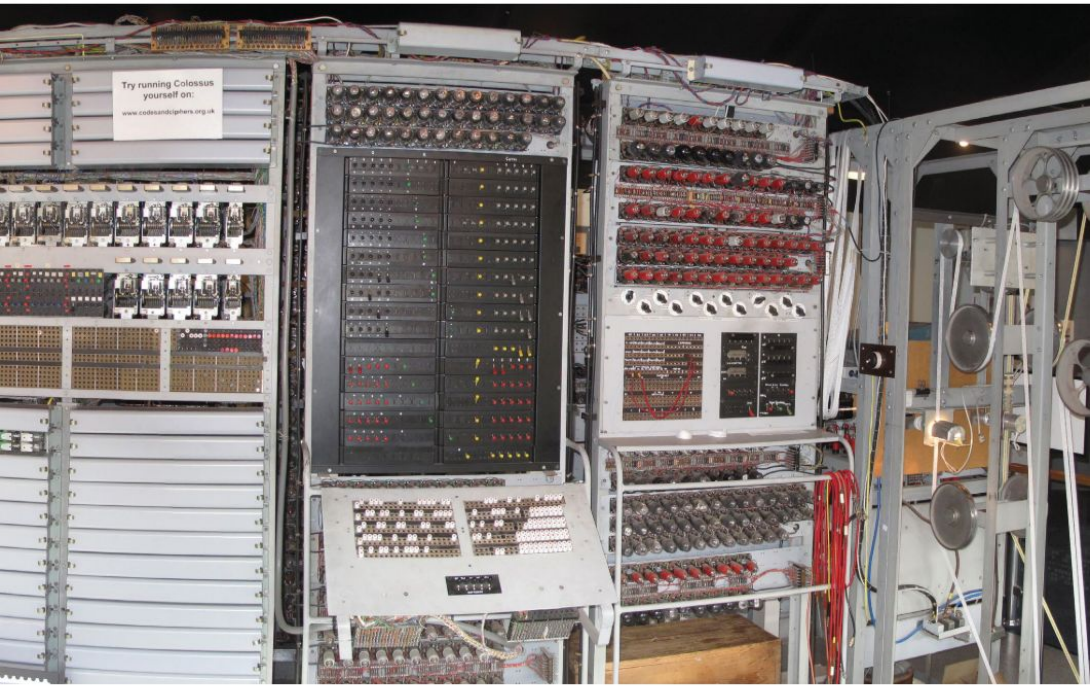

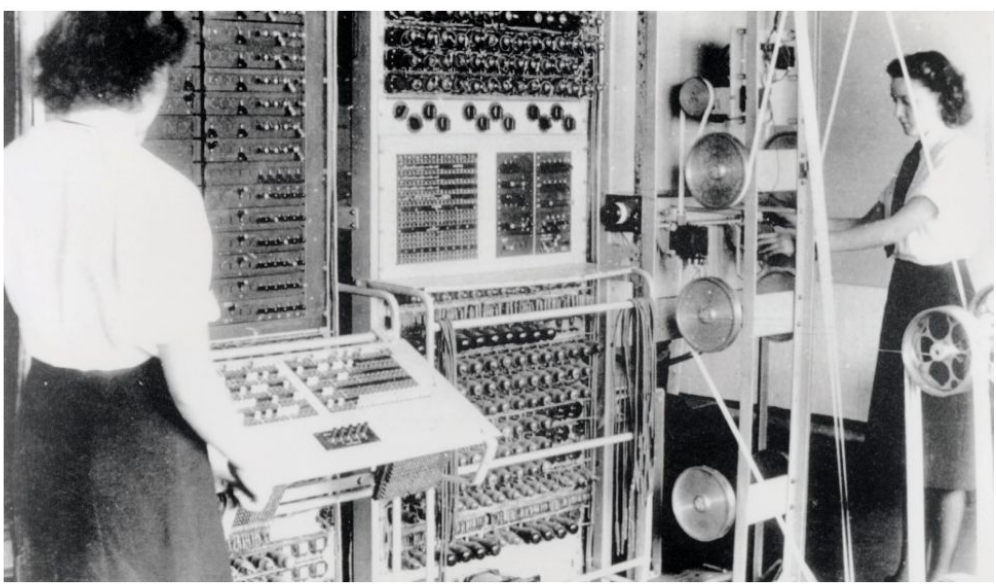

© TED COLES/WIKIMEDIA.

A Colossus Mark II computer being operated by Dorothy Du Boisson (left) and Elsie Booker (right), 1943.

CODE VS. CIPHER

Whether something is a code or a cipher is a question that mainly vexes magazine editors looking at almost 4,000 words of copy in which the terms are used frequently, but the two concepts are not interchangeable.

One way of looking at it is that codes change words, but ciphers change individual letters, or bits. In the context of Lorenz, the Baudot code took the words that were enciphered letter by letter by letter by the *Geheimschreiber*, and turned them into something else.

Another approach is that ciphers are for keeping things secret, whereas codes can be publicly known. An example here is Enigma, which was meant to be secret until Bletchley Park got its hands on it, and which used Morse code to transmit gibberish to other Enigma operators, who then transformed it back into readable German through knowledge of Morse and the secret Enigma machine.

Roman dictator Julius Caesar, though not blessed with a wireless network through which to propagate his orders, made use of a simple substitution cipher in which letters are shifted a set number of places in the alphabet. So, with a shift of five, the common phrase "Maximum PC, Minimum BS" becomes "Rfcnrzr UH, Rnsnrzr GX." You can see here an immediate weakness of a single-alphabet cipher: The words "Maximum" and "Minimum" look very similar even when enciphered. A polyalphabetic cipher, such as Enigma or Lorenz, fixes this, but Enigma had a cryptographic weakness—no letter could ever be enciphered as itself—that Lorenz didn't.

The Enigma machine resembled a typewriter.

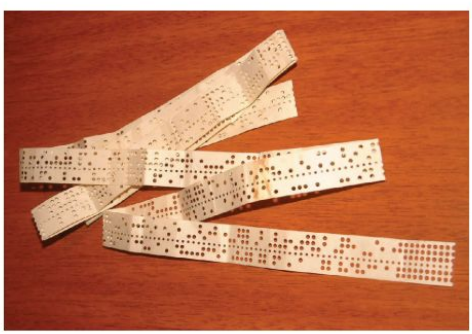

Paper tape with the punch holes of 5-bit Baudot code, and sprocket holes running down the middle.

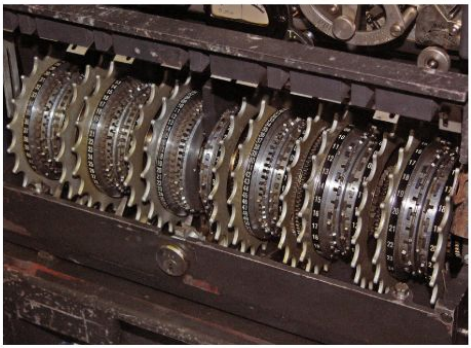

The wheels in the Lorenz machine each have a series of cams that can raise or lower, which were the settings Colossus calculated.

during the 1930s, and transferred to Britain at the outbreak of war), and a series of electromechanical devices known as Bombes, which simulated the rotors of the Enigma machine and could sift through possible settings to crack that day's wheel order and plugboard connections—effectively a brute force attack but using human brainpower, some wartime heroics (such as snatching a codebook from a sinking submarine), and user errors, such as transmitting two identical messages with different settings, to lower the number of permutations necessary. Lorenz, being far more complex, required a different approach.

At this point, Thomas H. Flowers, known as Tommy, the top man in research and development at the General Post Office Research Station at Dollis Hill, London, became involved. He was the son of a bricklayer, in contrast to the middle- or upper-class types usually found at Bletchley Park, and had earned a degree in electrical engineering by taking evening classes. He'd been exploring the use of electronics in telephone exchanges—usually powered by young women before this point in time—and was convinced he could build an all-electric one.

Which is how he came to the attention of Alan Turing, who asked the Post Office for help. "The issue started really with the Enigma when the German navy introduced a fourth rotor wheel," says Gannon. "A fourth rotor was just too big a job for an electromechanical device to sort out, so they set up two or three different projects." His first project with Turing was abandoned, but Turing was impressed enough to introduce Flowers to Max Newman, who was leading the assault on Lorenz in his "Newmanry"—the part of Bletchley Park working on automation.

"Turing wasn't directly involved in the whole Tunny question," says Gannon. "But he was a general consultant on all sorts of things at Bletchley Park, so got a little bit involved well before they even came up with the idea of Colossus. When they were stuck on how to decipher the messages that they were picking up, he suggested using a technique known as the Delta."

The math involved is too much for a humble PC magazine, but it involves repeated characters (a feature of German, which is full of SS, EE, TT, and LL). The ciphertext can be manipulated in such a way that certain relationships appear between characters that wouldn't be there if it was random, and from this the movements of one set of wheels in the Lorenz machine could be deduced. It was called Delta for the Greek letter, and the word "difference," because XOR is the same as modulo 2 subtraction.

"This became a key factor in what Colossus did," says Gannon, "as the way the Lorenz machine was built meant that

© WIKIMEDIA, INTERNATIONAL SPY MUSEUM, RICARDO FERREIRA DE OLIVEIRA

From left to right: **Tommy H. Flowers, telephone exchange expert and builder of the first electronic computer; Max Newman, British mathematician, codebreaker, and computer builder; C. E. Wynn-Williams, inventor of Colossus's counting device.**

XOR TABLE

INPUT		OUTPUT
A	B	A⊕B
•	•	•
•	x	x
x	•	x
x	x	•

A table showing the effect of the XOR function on 1 and 0 inputs.

two of the wheels generated an intermittent stutter in order to try to increase the periodicity of the wheels. By using the Delta technique, you could observe these stutters because you'd get a regularity, lines of spaces coming through. You could then statistically see through different layers of the cipher."

This method of attacking Lorenz was carried out in the "Testery," the part of Bletchley Park run by Ralph Tester, and supplemented by machines in the Newmanry. Flowers' first design for Newman was called the Heath Robinson (named for a British illustrator not dissimilar to Rube Goldberg), produced in conjunction with his colleague Frank Morell at the Post Office Research Station. Begun in January 1943, it was delivered to Bletchley Park in June, and worked pretty much straight away.

To understand how, we must step back a bit. In 1941, a Nazi radio operator made a mistake. He transmitted a message of 4,000 characters from Athens to Vienna using Lorenz, but it wasn't received properly. Vienna asked Athens to repeat the message, but didn't encrypt this request. The British immediately began to listen very carefully. The 4,000 characters were sent again, but without changing the key settings. This was very bad practice—in fact it, was forbidden—and to make things worse, the two messages weren't quite identical; the operator used abbreviations to make it shorter. In cryptography, this is known as a "depth," and can be used to break into the cipher.

Enter the enormous brains of Bill Tutte, a British/Canadian mathematician, and John Tiltman, a British Army intelligence officer. Tiltman deduced that it was a Vernam cipher using the XOR function, and was able to extract the key. This huge feat of intellect relied on Turing's observation that, once you had recognized and removed the effect of the first five wheels in the machine from the message, what was left, though still unreadable, was not statistically random. It took two months, but Tutte, backed by other members of the Research Section, was able to work out the complete structure of the Lorenz machine, despite never having seen one, from this one mistake by a radio operator in Athens.

It was this break-in that first Delta and then Heath Robinson were designed to exploit—removing the effect of the first five wheels from the enciphered message, so that manual methods

© WIKIMEDIA

WHAT WAS EVERYONE ELSE DOING?

Of course, the Germans weren't the only player in 20th-century geopolitics to encrypt their messages. The British took Enigma, added some extra bits to it, and produced Typex, a seven-rotor machine that the German *Beobachtungsdienst* (Observation Service) of Naval Intelligence never broke into despite reading most other British radio traffic during the war.

***Beobachtungsdienst* was based in Berlin, however, and Allied bombing raids destroyed many of its records. This, along with the British taking up the practice of superencipherment (enciphering a message using one method, then again using another) meant that by 1944, no British ciphers in use could be read.**

In the United States, the Japanese diplomatic cipher known as Purple was of more interest, and the intelligence it provided was code-named Magic. Purple was again a modified Enigma, but used stepping switches instead of rotors, and was known in Japan as the System 97 Typewriter for European Characters. The US Army Signal Intelligence Service reverse-engineered the system and had a machine that could read Purple by late 1940, but would not obtain the fragmentary remains of a genuine machine until 1945. As the Japanese ambassador in Berlin tended to brief his government in great detail using Purple (reviews of German assessments of the military situation, reports on direct inspections by the ambassador, including Normandy beach defenses, reviews of strategy and intentions, and reports of interviews with Hitler), Magic was of great interest to the British, and two of the reverse-engineered Purple machines ended up in British hands.

The Japanese naval cipher, JN-25, was trickier. Tiltman and Turing at Bletchley Park worked out quite early that it enciphered words as five-digit numbers using a codebook, then used a second codebook as superencipherment. Without the codebooks, however, they couldn't read it, and in 1940 the Japanese Navy wasn't sending many messages. But after Pearl Harbor, there was more traffic, and the Japanese habit of starting messages with "I have the honor to inform your excellency..." led to it being broken.

Two women at work at the US Army Signal Intelligence Service during WWII. The woman on the right is inputting ciphertext into an analog of the Japanese Type B Cipher Machine.

could do the rest. It used two synchronized paper tapes for data storage (each punched with up to five holes, one for every bit of the Baudot code), and was largely electro-mechanical apart from a few thermionic valves—key components of electronics for the first half of the 20th century until the invention of the transistor. The characters of the message were punched into the paper tape, which was passed at speeds of up to 2,000 characters a second past a photoelectric reader. Flowers' valve-based "combining unit" then applied the XOR logic that Tutte had figured out, passing the result to a counting unit designed by Dr. C. E. Wynn-Williams from the Telecommunications Research Establishment at Malvern. This used a thyratron—a sort of high-voltage thermionic valve—to count the number of 0s generated. The higher the 0 count, the more likely the key sequence was correct.

The Heath Robinson was slow and a bit unreliable. Paper tapes could stretch, putting them out of sync, and the display of the 0 count was overwritten with each new count, leading to lost data if not written down in time. Many of the shortcomings were dealt with as the machine was improved, but Flowers realized he could produce a machine that generated the key stream electronically. This meant a lot more valves. The idea that up to 2,000 valves could work together reliably was not popular. They weren't seen as reliable, with a tendency to blow like old-fashioned light bulbs.

"There was a dispute between Gordon Welchman [Assistant Director for Mechanization at Bletchley Park in 1943] and Flowers,

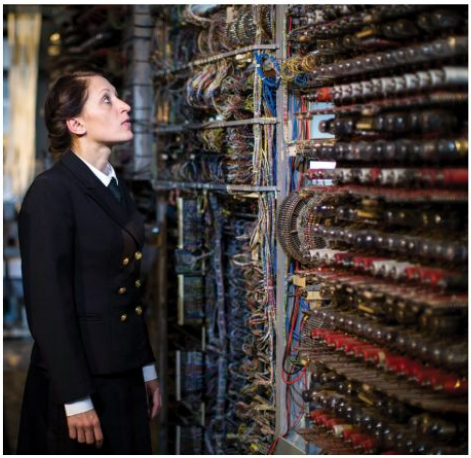

A Colossus computer, used during World War II at Bletchley Park, on display at The National Museum of Computing in the UK.

WHAT HAPPENED NEXT?

BILL TUTTE became a Fellow of the Royal Society of Canada in 1958 and of the London Royal Society in 1987. He became an Officer of the Order of Canada in 2001, and has a road named for him in Waterloo, Ontario. He died in 2002.

JOHN TILTMAN was the first non-US citizen to be inducted into the NSA Hall of Fame. He was awarded the British honors OBE, CBE, and Most Distinguished Order of Saint Michael and Saint George, as well as the Legion of Merit from the US Armed Forces. A consultant for the NSA until 1980, he died in 1982, in Hawaii.

C. E. WYNN-WILLIAMS who had studied with Ernest Rutherford before his wartime work, returned to Imperial College London, where he became Reader in Physics. In 1957, he received the Physical Society's Duddell medal in recognition of his work on the counting device that would be used in Colossus. UK Government Scientific Intelligence advisor, Professor R. V. Jones, wrote in 1981: "The modern computer is only possible because of an invention made by a physicist, C. E. Wynn-Williams, which may prove to be one of the most influential of all inventions." He died in 1979.

MAX NEWMAN would go on to help produce the world's first stored-program electronic computer in 1948 in the Royal Society Computing Machine Laboratory at the University of Manchester. Known as the Manchester Baby, it was a direct ancestor of the Ferranti Mark 1, which in 1951 was the first commercially available general-purpose computer. It was co-developed by Conway Berners-Lee, father of World Wide Web inventor Tim Berners-Lee. Newman turned down an OBE in 1946, but was already a Fellow of the Royal Society before the war, and received its Sylvester Medal in 1958. He died in 1984.

ALAN TURING was famously persecuted by the British government for his homosexuality, taking his own life as a result, and eventually being pardoned in 2013. After the war, he worked for the National Physical Laboratory, where he developed ACE, the Automatic Computing Engine. Its 25 kilobytes of memory, accessed at a speed of 1MHz, was described as "very large indeed." He moved to the University of Manchester to join Newman, where he consulted on the software for the Ferranti Mark 1, before becoming interested in mathematical biology and morphogenesis, the development of patterns and shapes in biological organisms, before the structure and role of DNA were known. He was made an OBE in 1946, and a Fellow of the Royal Society in 1951. A road and bridge in Manchester are named for him, and there is a statue of him in the city, too. In 1999, *Time* magazine named Turing as one of the "100 Most Important People of the Century." He died in 1954.

GORDON WELCHMAN moved to the United States in 1948, where he taught the first computer programming course at the Massachusetts Institute of Technology. Among his students was Frank Heart, who would go on to help design the first routing computer for the ARPANET, the predecessor of the Internet. Welchman would work for Ferranti, and on secure communications systems for the US military at Mitre. He wrote a book about his wartime work in 1982, leading to him losing his security clearance and thus his position as a government consultant. He died in 1985.

© JACK TAYLOR/GETTY IMAGES

where Welchman was trying to get Flowers kicked out of Bletchley Park because he was wanting to use too many electronics," says Gannon. "His ideas would have used up all the electronic valves in the country." Gannon's theory, based on his research, is that Turing intervened and took Flowers' side. This was just as well, as Flowers' work at the Post Office had shown that valves blew from stresses that were greatest when they were powered up. His solution? Don't turn the machine off. A team of 50 people set about building something completely new.

Colossus Mark I, with 1,600 valves, arrived at Bletchley Park on January 18, 1944, and attacked its first message on February 5. A Mark II machine, with 2,400 valves, became operational on June 1, just in time for D-Day, the Allied invasion of France, and would be joined by 11 more. The Mark II had 12 rings of thyratrons that simulated the 12 wheels of the Lorenz machine, and a six-character shift register of Flowers' design which, along with the five electronic counters, gave it five-way parallelism, so it was five times faster than the original and easier to operate. Its speed was still limited by the paper tape, but it meant the British were reading high-level Nazi communications, including those signed by Hitler, just hours after they were sent.

Colossus was built for one task, and so couldn't be called a general-purpose computer. Its programmability was limited, with the programs held in the positions of switches and jack panel connections, but its algorithms could cover five billion combinations of variables. Of course, the Germans hadn't sat still, and by August 1944, the wheel settings on all Lorenz machines were being changed daily, leading to a lot of work.

After the war, Colossus was a secret for 30 years. Units 1–10 were dismantled, their parts returned to the Post Office or Newman's Royal Society Computing Machine Laboratory at the University of Manchester. Flowers was ordered to burn all documentation. "That was a terrible mistake," he said later. "I was instructed to destroy all the records, which I did. I took all the drawings and the plans and all the information about Colossus on paper and put it in the boiler fire. And saw it burn."

Units 11 and 12 survived, however, and as the Government Code and Cipher School was renamed Government Communications Headquarters (GCHQ) and moved to a new building at Cheltenham, they, plus two replica Lorenz machines, were taken along. They remained there until 1960, used mainly for training and for checking that one-time pads (an unbreakable form of sending secret messages) were sufficiently random.

Those who knew about the machines, however, dispersed into academia and industry. The idea that high-speed digital computing was feasible and reliable was a powerful one, but it didn't catch on everywhere; when Flowers applied for a bank loan to build another machine, he was rejected because the idea was not considered likely to work. Constrained by the Official Secrets Act, he could not argue that he'd already built one. He remained at the Post Office Research Station, where he became head of the switching division, finally building his all-electric telephone exchange in 1950. He moved to Standard Telephones and Cables Ltd, and retired in 1969. He published a book on telephone exchanges in 1976, received an honorary doctorate in 1977, the Martlesham Medal in recognition of his achievements in computing in 1980, and in 1993 was awarded a college certificate, having completed a course in basic information processing on a PC. He died in 1998.

Recognition was not something the man who built the first computer sought in his lifetime, but it has come to him, slowly. There are at least two roads named for him, and a computing center for young people, The Tommy Flowers Centre, opened in 2010 near where he was born. At the R&D center operated by British Telecom (BT)—the successor to the Post Office in running Britain's phone network—a life-size bronze bust of Flowers was unveiled in 2013, and in 2016, BT opened the Tommy Flowers Institute, a training center for postgraduates moving into the telecommunications industry.

And despite the blueprints being destroyed, enough material survived in engineers' notebooks (a surprising amount of the material had made its way to the US) for a Colossus Mark II to be rebuilt between 1993 and 2008. The optical tape reader was rebuilt by its original designer, Dr. Arnold Lynch, and the working unit stands today in the position occupied by Colossus No. 9 at Bletchley Park, now The National Museum of Computing. It's the size of a bus, and makes a noise like the world's worst industrial metal band falling down the stairs.

A competition to celebrate the rebuilt Colossus pitted it against modern computers in receiving and decoding Lorenz messages. The winner, Joachim Schüth, using a 1.4GHz laptop and his own program written in Ada, took 46 seconds to find the settings for all 12 wheels. Colossus took three hours and 15 minutes. "My laptop digested ciphertext at a speed of 1.2 million characters per second, 240 times faster than Colossus," said Schüth at the time. "If you scale the CPU frequency by that factor, you get an equivalent clock of 5.8MHz for Colossus. That is a remarkable speed for a computer built in 1944." Remarkable indeed. Intel wouldn't get there for another 35 years.

KONRAD ZUSE

Working alone, with financial help from the government, German engineer Konrad Zuse produced the first programmable computer in... 1941? Zuse's work is often overlooked, but he began work in 1935 after becoming bored with the calculations he had to do by hand in his job at an aircraft factory. In 1936, he produced the Z1, a floating-point binary mechanical calculator with limited programmability, and also submitted some interesting patents. His work was destroyed in an air raid, but in 1939 he was called up for military service, and produced the Z2, an updated Z1 built with telephone relays.

While not a member of the Nazi Party, Zuse's work benefitted the regime and contributed to its missile programs. He was advised to work with vacuum tubes, but thought the idea "crazy." More bombing destroyed the Z3, a binary 22-bit floating-point calculator featuring programmability with loops, but the Z4—a 32-bit floating-point programmable computer built from mechanical and electromechanical logic—was missed by the bombers.

Post-war, IBM took an option on his patents, providing an income, and by 1950 had delivered the Z4, the only working computer in continental Europe, to a Swiss university. The Z series would go up to Z43; the Z22 was the first computer to use magnetic storage. He submitted a PhD thesis to the University of Augsburg, but forgot to pay, so was rejected. He continued to work on programming languages, large format plotters, and even published a book on the digital structure of the universe. A painter in his spare time, Zuse died in 1995. The Zuse Institute for applied mathematics and computer science in Berlin is named in his honor.

© WOLFGANG HUNSCHER, DORTMUND / WIKIMEDIA.

EXPLORE THE BEST LINUX & OPEN SOURCE SOFTWARE

From the best distros to system security, utilities and fun applications, we've rounded up everything you could ever need

Linux & Open Source Annual 2021
Everything you need to master open source software and operating systems
180 pages of expert tips & tricks
ON SALE NOW
100% UNOFFICIAL
ISBN 978-1-80023-059-0
9 781800 230590

FUTURE

Ordering is easy. Go online at:

www.magazinesdirect.com

Or get it from selected supermarkets, newsagents, or bookstores

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

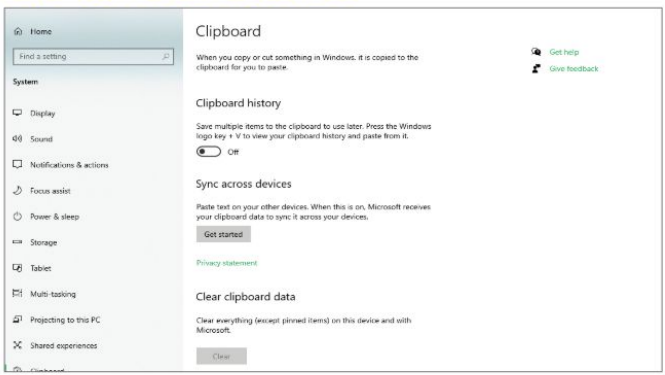
Home
Find a setting
System
Display
Sound
Notifications & actions
Focus assist
Power & sleep
Storage
Tablet
Multi-tasking
Projecting to this PC
Shared experiences
Clipboard
When you copy or cut something in Windows, it is copied to the clipboard for you to paste.
Get help
Give feedback
Clipboard history
Save multiple items to the clipboard to use later. Press the Windows logo key + V to view your clipboard history and paste from it.
Off
Sync across devices
Paste text on your other devices. When this is on, Microsoft receives your clipboard data to sync it across your devices.
Get started
Privacy statement
Clear clipboard data
Clear everything (except pinned items) on this device and with Microsoft.
Clear

CLOUD CLIPBOARD

Who knew copying and pasting could be such fun? A little sarcasm there, but if you have a PC and a laptop, for instance, this could boost your productivity. Head to cloud settings in Win 10, sync across devices, and follow the instructions. It's a simple tip but handy if you find something you want to copy from laptop to desktop, or vice versa. Press Win-V to view your clipboard, and turn on clipboard history to access old content.

MAKE – USE – CREATE

60 Run and manage your Docker containers easily

64 Crack the Enigma code on your Raspberry Pi

68 Build your own ITX VR-ready gaming PC

© MICROSOFT, DOCKER, RASPBERRY PI FOUNDATION

SAM LEWIS
STAFF WRITER

HOME OFFICE

They say you should keep business and pleasure separate—not in this industry. I love working from home; being next to the kitchen certainly has its perks! Going into the office with the team will have its advantages, of course, and I'm looking forward to it, but my home setup is getting pretty good. I've just switched from the 34-inch MSI PS341WU monitor to a 37.5-inch HP Z38c curved display. At this rate, I'll be using a projector by Christmas!

With some of the testing I have been doing, paired with a very capable PC, I've been able to revisit some games of yesteryear the way they were intended to be played. I've sunk myself into *Alien Isolation* once again. It's a game I shouldn't enjoy, considering it scares me to the nth degree, yet I always find myself drawn back to it. With the curved monitor, the immersion has only greatened—which has its pros and cons....

As an ex-PlayStation user, I've also now been able to enjoy some of the Xbox hits that I swore never to touch as a narrow-minded teen, *Forza* being one of them. Playing in ultra on a 21:9 widescreen is where this game thrives, and I can see the appeal. PC gaming is something I've always wanted to dive into and I can do that now. With VR testing and gaming peripherals heading my way, I'm excited to explore this further.

↘ submit your How To project idea to: comments@maximumpc.com

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products. To learn more, visit www.ifixit.com.

Two screws secure the Lightning port, making this high-wear part easy to replace if it breaks.

It's fully optimized for the iPhone, but with some fiddling around, we're sure a USB-C port could be swapped for the Lightning adapter.

Backbone One

Cracking open the casing reveals a rainbow road of cables—we hear you, *Mario Kart* fans.

The use of Phillips screws rather than glue to hold it together makes the Backbone One easier to fix.

The left side has an almost identical setup to the right side—all the way down to the Phillips screws and ZIF connector.

BACKGROUND

The Backbone One sneaks into the mobile gaming market with a handful of interesting features, including a pass-through Lightning port, headphone jack, and some dedicated controls, such as gameplay recording and instant muting. But is this compact iPhone controller as repairable as it is tactile? Let's see!

MAJOR TECH SPECS

- D-pad and thumbstick on the left
- A-B-X-Y buttons and thumbstick on the right
- Shoulder and trigger buttons on each side
- Lightning port for pass-through charging and accessories
- 3.5mm headphone jack
- A soft-touch Smart Case and a Lightning-to-USB-C cable

KEY FINDINGS

- This Joy-Con-esque controller may look like any other mobile gamepad—but it packs a few surprises. A couple of interesting omissions, though, are Bluetooth and an internal battery. The Backbone One connects to and draws power from your iPhone.
- We're off to a fast start thanks to Phillips screws on the right grip. Beneath the plastic shell, some colorful cables—and our first look at the internals. We free the interconnect cable from its ZIF connector claw and unplug the two Lightning connector cables from the motherboard. A few more Phillips screws later, we can extract the shoulder button, trigger, and our first PCB.
- Let's see what this board brings to the game: joystick (soldered to the board—not a great start), six clickable contacts (tactile dome) for the A-B-X-Y buttons, orange Backbone App button, and option button. It also has Skyworks SKY52101-11 IoT SoC, Avnera AV3425C single chip analog SoC for smart Lightning headsets, and soldered-on bumper button, Hall-effect sensor for the trigger button, and Lightning port to pass through power to your phone or sound to your headset.
- Over on the left side we find an almost identical setup, but we're momentarily stumped trying to remove the trigger button. Turns out, it's secured by just one well-hidden screw. The left board offers another soldered joystick (like the Razer Kishi), four gold contacts for the direction pad and four holes to keep its silicone cover in place, and two clickable silver contacts (tactile dome) for the screenshot/record and option buttons. You can also find a soldered 3.5 mm headphone jack, bumper button, and Hall sensor for the trigger button.
- The sliding clamp mechanism uses a pair of flat spiral springs that unroll when you pull the handles apart. This provides more consistent compression and a tighter grip than the Razer Kishi, which uses helical extension springs to embrace your phone.
- Two screws secure the Lightning port. With the right tools, some case-modding, and rewiring, you might even be able to swap it out for a wider USB-C port to make it Android-compatible. The interconnect cable wiggles in the casing but can't be removed without causing irreparable damage. We free the springs with our Minnow Precision Bit Set and save the game before the final boss. We have beaten the Backbone One.
- Repairability Score: 5 out of 10 (10 is easiest to repair). Disassembly is straightforward with components held in place with screws. The Lightning connector can be easily replaced by removing the cover and two screws. No adhesive is used, but some screws are difficult to reach. Shoulder buttons are soldered to their boards, requiring micro-soldering skills for replacement. Both joysticks are soldered to their boards, as is the pass-through Lightning port.

How to Use Photoshop Part 3

YOU'LL NEED THIS

ADOBE PHOTOSHOP CC 2018 ONWARD
To view compatibility visit: https://helpx.adobe.com/photoshop/system-requirements.html.

WINDOWS 10
64-bit, version 1809 or later.

LOOKING TO ADD A LITTLE WOW factor to your images or create abstract cover art for a project? Adding a glow effect is a great way to achieve an eye-catching look. It can transform standard-looking imagery into unique art that can serve many purposes. We're focusing on using the selection tools once again and refining edges to get a perfect perimeter. Refining an edge is a great way to select exactly what you want; sometimes AI can't quite hack it.

Other areas we're focusing on are blur and hue/saturation effects, and adding them together to create a nice even glow and an aesthetically pleasing end product. These effects enable you to create depth in an image, putting more emphasis on the area you want to glow. The part of the image you choose to glow is important—it needs to be at the top of the hierarchy, because it's what people will look at first. Hopefully, by adding these tutorials to your artistic arsenal and putting your imagination to the test, you can create something visually grabbing. **–SAM LEWIS**

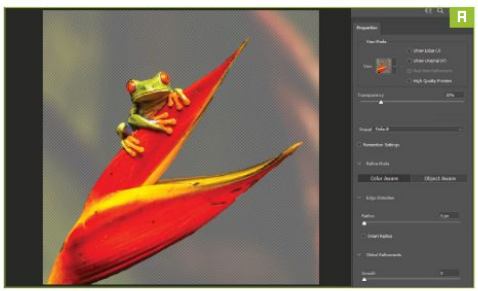

1 PREPARE YOUR IMAGE

Like every tutorial so far, you need an image; try to find one with a subject in the foreground. We're using a nature shot to showcase how to add a little sparkle to an image. Begin the same way as before: Load your image into Photoshop and make an original copy (Ctrl-J) to duplicate a layer. Rename the bottom layer "original [insert name]," then lock it, and hide it away. We've chosen a very interesting-looking frog for this example, and the end goal is to make its eyes glow.

›› Using the selection tool, or any other selection method of your choice, highlight around the subject of your image. For ours, we have highlighted the frog and the plant, because they are both in focus.

›› Once selected, click "Select and mask..." at the top of the screen. This opens a new view with refining tools. In this section, we're refining the edge of our selection. Different view options are available on the right to help differentiate the foreground from the background.

›› We've selected the onion skin view [**Image A**]; you can slide the transparency up or down to make your selection more or less prominent. Once you're satisfied that you can see your selection clearly, head over to the left of the screen, and click the Refine brush tool. This should be the second icon down in the left-hand tool panel. Draw around rough areas of your selection—you can use the sliders in the panel on the right to adjust the sharpness of the brush. Once you have finished refining your edge, you must change the output to "New layer with layer mask" (you can find this at the bottom-right of the right-hand panel). Hit "OK."

2 GROUP WORK

Next, we are going to select the glow part of our selection. Hold Ctrl and click the mask layer you just created. This should select your mask selection. Then, using the selection tool, highlight the areas you want to glow, either adding or subtracting areas using the Alt key. With your selection, click the image thumbnail in the layer instead of the masked section, and duplicate the layer: Ctrl-J. We then put the new layer into a group—with the layer selected, hit Ctrl-G, then hide the group [**Image B**].

3 CONTRASTING VIEWS

Moving on, we need to make the foreground stand out further. To do this, click your first layer (the unedited picture, not the backup locked layer), click the circle icon at the bottom-right called the adjustment layer, used in the last tutorial, then hit "Color look up."

›› This opens a new Properties tab; click in the 3D LUT file section and find the moonlight effect. This darkens the background, creating more emphasis on the subject.

›› After increasing the contrast between foreground and background, unhide the group created earlier, right-click the "glow" layer, and change it to a Smart Object. Duplicate this, and on the duplicated layer, change the layer type from "Normal" to "Linear dodge (add)." To smooth the edges, go to "Filter → Blur → Gaussian blur," and change the radius to around five pixels. Then duplicate once more.

4 BLURRED VISION

Once that is done, we need to increase the glow. Click the latest duplicated layers—they should all be in the group created earlier. Click into the Gaussian blur filter in the layer and change the radius to 100. This adds a wider glow. Duplicate this again: Ctrl-J. Learning these keyboards shortcuts is incredibly useful for increasing efficiency when creating projects. Follow the same step again, click into the Gaussian blur filter in the layer, and change the radius this time to 250, further increasing the glow. Repeat once more, now changing the Gaussian blur radius to 500.

›› Then click "Group layer → Adjustment layer → Hue/saturation." This opens a hue/saturation layer in the group [**Image C**]. Within this, click the left thumbnail in the layer, and a Properties tab now appears. At the bottom of the Properties tab is an "Adjustment clip icon"—it looks like a square with an arrow coming out of the right and pointing down. Hit this to apply. In the tab, click "Colorize," increase the saturation, and then select what color you want your glow to be using the slider. We decided to go with a strong yellow and increased the brightness to over-exaggerate the effect.

5 GLOW WITH THE FLOW

To create more of a contrast from the glow and the rest of the foreground, we need to use another hue/saturation adjustment layer. Click the layer of the original image and the mask. Then add the adjustment layer, just as we did earlier, and be sure to clip it, too. Drop the saturation a little and do the same with the brightness—this should make the other areas of the foreground darker compared to the glow.

›› To finalize, we dropped the opacity of the group layer to 90%, as we felt the glow was too strong and wanted a little of the original image to come through [**Image D**]. You can always go back and alter things slightly. Use these tutorials as guides; each image will be different and art is subjective, so choose what you feel works best. Try using this technique to create interesting imagery for cover photos, or to enhance certain parts of pictures, maybe even album covers—who knows?

© ADOBE

YOUR TOOL BOX

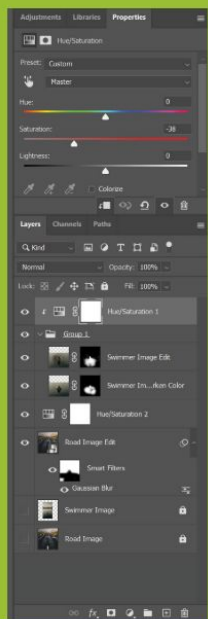

Using the Refine Edge tool can take your work to the next level. Making a selection can be a pain; as good as Photoshop's AI is, it doesn't always select what you want, or it can leave sharp edges along the edge of a cutout. An example of when the Refine Edge tool can help is when making a selection around hair. If you ever need to do this, patience is certainly virtue, so take your time!

The tool is used for complex selections in Photoshop. To get the best results, make a semi-rough selection with a tool of your choice, then click "Select and mask" to open the Refine Edge tool window. On the right side, you see the Properties tab.

The View mode is great for different types of imagery; it changes the background and colors, so you can see the outline of your selection easier. Beneath that is the Transparency setting for the view—alter this to your liking. The next setting of interest is the Refine mode—there are two options: Color Aware and Object Aware. Try both to see what gives the best results. The Radius slider is used to increase or decrease the radius of the edge detection. With all these options, you can preview as you go, so keep tweaking until you get what you need. The global refinements are great for getting the edge you need. All generally do what you'd expect. Practice with mock images to understand how these tools affect the edge of a selection. Finally, at the bottom are the Output settings. You can output a selection, layer mask, new layer, new layer with layer mask, new document, or a new document with a layer mask.

On the left are the tools. You can use all of these to add or subtract from the image, holding the Alt key to differentiate. At the top is the Selection tool, then the Refine edge tool. This is the one you will probably use most to follow around the edge of your selection. Under that is a Brush tool to edit by eye. Beneath those are the Object Selection tool and the Lasso tool, handy for selecting specific parts of the image if you aren't getting your desired edge. Finally, there are two standard tools for navigation: the hand and magnifier.

With patience and practice using all the properties and tools, you will be able to make a clean selection out of any image. If the selection is precise and tidy, it helps with blending and making things look more natural.

Manage Your Docker Containers Easily

YOU'LL NEED THIS

A PC OR MAC RUNNING DOCKER
Better still, an Ubuntu-powered PC running Docker, Cockpit, and Cockpit-Docker).

LAST ISSUE WE REVEALED how to set up your own headless server running Ubuntu Server, complete with remote access via the Cockpit web-based UI. We ended by introducing you to the world of Docker and containers, which enable you to run individual applications and services within their own isolated environments for security and stability reasons.

This tutorial picks up pretty much where that feature left off—so you need to refer back if you've not yet installed Cockpit, Docker, or the Cockpit-Docker plugin. If you're looking to get started with Docker on another platform, the good news is that you can run it on Windows, macOS, and other flavors of Linux, too, complete with your choice of user-friendly front end—the box on page 62 reveals some of the options available.

We'll open by revealing how Linux users can run containers as a non-root user for security purposes, plus step you through the process of finding, downloading, setting up, and running containers on your server. We'll even show you how to get around any missing features in Cockpit-Docker by bypassing it and using the Terminal in conjunction with your PC's text editor to quickly get more complex containers up and running. Ready to transform your new server? Let's get started! –NICK PEERS

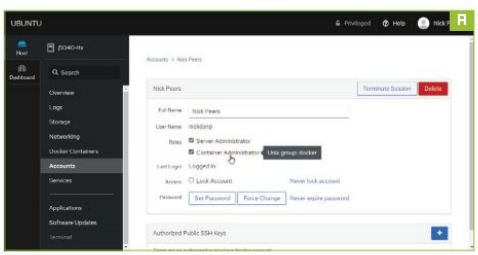

1 SET UP DEDICATED DOCKER USER

Once you've got Docker and Cockpit-Docker (or your choice of UI) up and running, Linux users should visit https://docs.docker.com/engine/install/linux-postinstall/#manage-docker-as-a-non-root-user for a guide to managing Docker without requiring access to "sudo." Those running Cockpit-Docker have no need to create the required "docker" group—it's been created for you. Switch to "Accounts" in Cockpit and click your user account. A new "Container Administrator" role has been added—checking this [**Image A**] adds your user to the "docker" group for the purposes of administering Docker from the command line, without having to precede commands with `sudo`.

›› Although Docker now supports rootless containers, which means they no longer need access to root for security purposes, they come with some restrictions that can make them impractical in all scenarios, specifically those containers that are accessed through any ports lower than 1024. Rootless Docker isn't set up by default, although it's simple to do in Ubuntu if you feel so obliged (see https://docs.docker.com/engine/security/rootless/).

›› An alternative workaround exists, which works on a container-by-container basis. This involves configuring each container to run using a specific user—including non-administrators—rather than the main root account. This entails creating a dedicated user for that very purpose. Switch to Terminal in Cockpit and issue the following command:

```
$ sudo adduser docker --ingroup docker
```

›› This creates a dedicated "docker" user and makes it part of the "docker" group—you're prompted to create a password during setup. We'll be running all containers in this tutorial through this user and group, and you need to identify the docker user (UID) and group (GID) IDs:

```
$ id docker
```

›› Make a note of these—our Ubuntu server's "docker" user, for example, returns a UID of 1001, and GID of 998.

2 SET UP CONFIGURATION FOLDERS

Container images are volatile storage, so when you shut down or restart a container, all changes are wiped out and the container starts up afresh. To prevent all your changes from being lost, Docker uses a system of mount points, which enables you to connect the container to folders on your hard drive, where data such as configuration settings can be stored securely.

›› The best place for configuration folders is a common folder within the "docker" user's "home" folder. Log out of Cockpit as your own user, then log back in as the "docker" user. Open the Terminal and create your folder:

```
$ mkdir containers
```

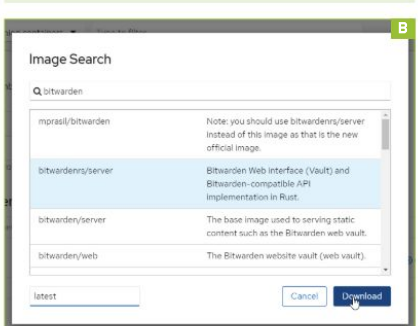
B
Image Search
bitwarden
mprasil/bitwarden
Note: you should use bitwardenrs/server instead of this image as that is the new official image.
bitwardenrs/server
Bitwarden Web interface (Vault) and Bitwarden-compatible API implementation in Rust.
bitwarden/server
The base image used to serving static content such as the Bitwarden web vault.
bitwarden/web
The Bitwarden website vault (web vault).
latest
Cancel
Download

© UBUNTU, DOCKER

MANAGING CONTAINERS

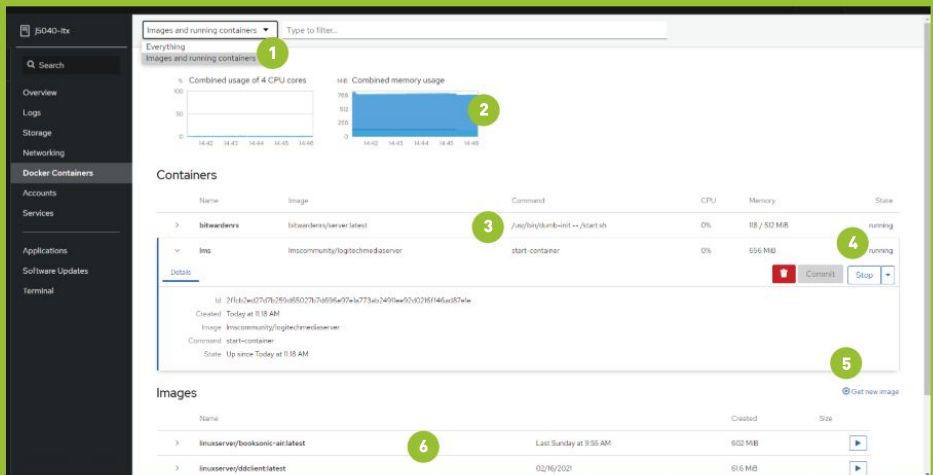

1. FILTER VIEW
Use the drop-down menu to switch between showing every container and only those containers that are currently running.

2. RESOURCE USAGE
Keep an eye on the CPU and RAM usage across all of your containers here. RAM usage is a little misleading, however, because it doesn't show your server's full RAM.

3. CONTAINER DETAILS
Any containers meeting the drop-down menu or filter's criteria are listed here, with at-a-glance stats helping you track resource hogs.

4. CONTAINER CONTROLS
Click ">" next to a container and you can access shortcuts to deleting, stopping, and restarting the container. Click the container to access its own status screen.

5. GET NEW IMAGE
Click here to search the Docker hub for images you can download for use as containers. Also use this to download the latest version of images you're already using.

6. IMAGE DETAILS
Click the play button next to an image to create a container from it; or click ">" to access details about the image, plus an option to delete it if no longer required.

›› By default, only the "docker" user has read/write access to this folder. Log out and back in as your own user account—if you've not already added this account to the "docker" group as outlined in the previous step, do so now (non-Cockpit users should type `sudo usermod -a -G docker username` to do so).

›› This gives your user account read-only access to the containers folder. If you needed read/write access for any reason, give the "docker" group full access from the Terminal:

```
$ sudo chmod -R 775 /home/docker/containers
```

3 DOWNLOAD DOCKER IMAGES

Time to start using Docker properly. Everything you need to run a container is housed in a Docker image, and there are thousands of examples at https://hub.docker.com, which is where Docker pulls the images from. Start your search for suitable containers here, using the box over the page for inspiration. You'll find many applications feature multiple times—if there's no obvious official entry, look for one from LinuxServer, or for the most popular or highly rated versions.

›› Once you've identified the image you wish to use, return to Cockpit and select "Docker Containers" from the left-hand menu. Click "Get new image" next to "Images," type the name of your chosen container into the Image Search window that pops up, and a list of matching containers appears.

›› Select your chosen container from the list and click "Download" [**Image B**]. After a short pause, you should see the image being pulled under "Images." Once complete, the image is available for you to create a container from.

4 SET UP YOUR DOCKER IMAGE

To create a container from your downloaded image, click the play button next to it to bring up the Run Image dialog. This contains all the configuration options you need to successfully run the container on your system. While some are automatically added by the image, it's still a good idea to open the corresponding web page on http://hub.docker.com to see what options are available to you.

›› Start by changing the Container Name from the existing title (which is a random "adjective_surname" pair) to something more recognizable. Leave the "Command" field as it is, then decide if you wish to limit the container's access to memory and CPU by checking the relevant box to set a RAM limit and/or alter the CPU priority—1,204 is the default setting here, which can be reduced to constrain the container or increased to give it more weight as seen fit [**Image C**]. Note, while RAM

WHAT, NO COCKPIT?

The beauty of Docker is that it works across Windows, Mac, and Linux, and you don't need Cockpit-Docker to provide yourself with a graphical front end. Windows and Mac users get such a front end by installing Docker Desktop (www.docker.com/products/docker-desktop), while Docker functionality—complete with web-based UI—is also available for QNAP (Container Station) and Synology (Docker) NAS drives.

There are plenty of alternatives, too—many of which also work with different flavors of Linux. If you're looking for a web-based UI, then Portainer Community Edition (www.portainer.io) is the tool of choice for many. It has all the options you'll find in Docker-Cockpit, and more besides. It's installed as a standalone Docker container:

```
$ docker volume create portainer_data
$ docker run -d -p 8000:8000 -p 9000:9000
--name=portainer --restart=always -v /var/run/docker.
sock:/var/run/docker.sock -v portainer_data:/data
portainer/portainer-ce
```

After installation, access it through your web browser: http://192.168.x.y:9000.

Whichever front end you choose, adapting this tutorial should be quite straightforward—indeed, Portainer opens up access to a wider range of tools from your web browser (and, yes, you can run it alongside Cockpit-Docker).

limits are absolute, CPU priority settings are only enforced when resources are scarce. Unlike the other settings on this screen, both can be amended later if you wish to make changes.

›› If you want to be able to access the container using tty, leave "With terminal" checked (if in doubt, uncheck it for security reasons). Leave "Link to another container" unchecked for a single stand-alone container.

5 CONFIGURE NETWORKING

By default, your container communicates with the outside world using a "bridge" network connection. This works by mapping ports on your host PC to those within the container itself, enabling you to communicate with the container through your host PC's IP address. Although Docker supports different types of networking relationships between container and host, Cockpit-Docker only works with the standard bridge configuration. The simple workaround here is to set up your container via the command line, as outlined in step 8.

›› The "Ports" section is naturally where you set up these mapped ports. Leave "Expose container ports" checked and you should find the image has already defined the container's ports that need access. The port on the container is listed on the left, and it's up to you to choose which host port to map it to. In most cases, a simple 1:1 relationship is fine (so port 3012 to port 3012), but on more commonly used ports (typically 80 or 443), you need to set a different port here [**Image D**]. If you need to define additional ports, simply click the "+" button next to an existing port to add a new one.

›› When choosing a different port to use, check out Packetlife.net's cheat sheet (https://packetlife.net/media/library/23/common-ports.pdf), which reveals which port numbers to steer clear of to avoid clashing with common services.

6 VOLUMES AND VARIABLES

The "Volumes" section is where you set up the mount points that connect virtual drives on your container to physical folders on your hard drive. There's usually a "/config" folder, which is where your settings are stored so they survive container restarts, but you may want to map further folders, too (such as access to a folder on your PC containing all your audiobooks). Here it's a case of manually entering the full host path (such as "/home/docker/containers/booksonic/config") and setting the folder to default, read-only, or read-write as required. Read-only ensures the target folder's contents can't be altered by the container.

›› The next section covers the environment variables your container requires to run—these are marked with the "-e" tag on the container image's home page on http://hub.docker.com. Again, many should be predefined, leaving you the simple task of verifying them and making changes where appropriate [**Image E**].

›› To run the container within your Docker account for security purposes, you need to create the following two keys: PUID and PGID. Their values should match the UID and GID values you recorded using the "id" command in step 2. If you're working with LinuxServer images, these keys should already be in place, ready to be edited.

›› There's one final setting to consider before you launch your container: "Restart Policy," which determines what happens when the container is stopped or halted (say, after a crash). There are four basic options: "No," "On Failure," "Always," and "Unless Stopped." They're largely self-explanatory—the key difference between "Always" and "Unless Stopped" is that "Always" leads to the container being restarted automatically after a server reboot, while "Unless Stopped" doesn't.

7 LAUNCH AND TEST

Once you've configured your container, it's testing time. Here's where containers prove their worth: You can't do any damage, so if you make a mistake, you can simply delete the container and start again. Go through the settings one final time, then click "Run." You're returned to the main Docker Containers screen,

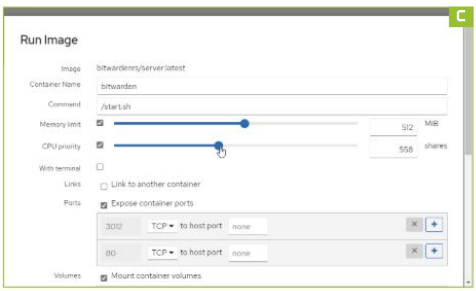

© UBUNTU, DOCKER

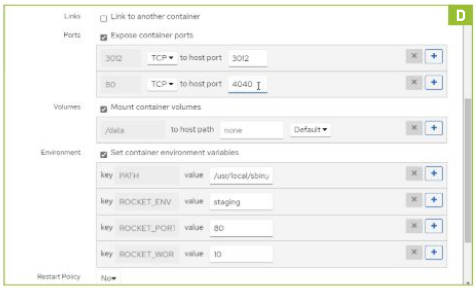

where you should see your newly created container now appears under "Containers," its "State" listed as "running."

›› Click this and you can view your container's status, plus access controls for starting, restarting, stopping, and deleting the container. Scroll down to reveal a console screen enabling you to see what's happening behind the scenes, plus a "Change resource limits" button, which lets you change RAM and CPU allocations without having to destroy and recreate the container.

›› If the container fails to start or the console reveals an error, simply shut it down and restart it. If the same problem occurs, shut down and delete the container, then try creating it afresh, double-checking your settings. Some containers may appear to exhibit errors while still working—once it's up and running, and an IP address is assigned (in bridged mode, this IP address is different from your host PC's IP address, which is the address you use to communicate with the server), you can check to see if you can connect from a client app or through your web browser. Once verified, you can go on to configure the server itself within its own web UI—when you shut down, restart, or recreate the container, all your changes are reapplied from your config folder.

›› One last thing: Most Docker images are periodically updated, but there's no option to automatically update them. Instead, you need to pull the image afresh.

8 RUN CONTAINERS FROM THE TERMINAL

Finally, we've seen how to set up and launch containers from Cockpit-Docker, but there are times when its "Run Image" controls aren't sufficient to meet your needs. It doesn't support passing commands such as `--device=/dev/dri:/dev/dri` (required to support Intel hardware acceleration in Jellyfin, for example) or using a different network relationship to the default bridge mode. It can also be time-consuming to set multiple things up again and again as you experiment with settings.

›› Thankfully, this is where a combination of Docker's CLI interface and Cockpit's support for full copy and paste in the

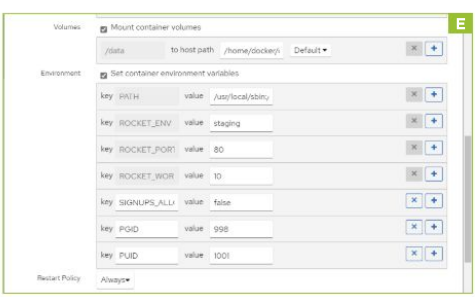

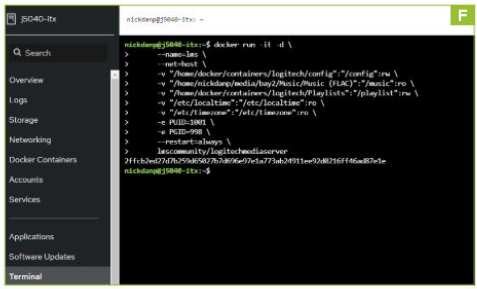

Terminal reap dividends. First, locate the image's page on https://hub.docker.com, where you should find a lengthy Docker CLI command with all the controls, environment variables, volumes, and so on listed on separate lines. Select all this text, then copy it to a blank text file in your text editor. Save the file for future reference, then make the changes required to the script in your text editor—setting the correct volume points, for example, or adding new commands as required (for example `--net=host` to connect the container directly to your server's host network). Make sure you add `-d` to the `docker run` line to ensure the container runs detached.

›› Once done, save your file, then press Ctrl-A to select all the text, switch to Terminal in Cockpit, and press Ctrl-V to paste it all in as shown above [Image F]. Press Enter and the container magically creates itself and starts up—and it appears in Cockpit-Docker, too, enabling you to monitor and control it from there.

YOUR FIRST THREE CONTAINERS

Containers are easy to create and destroy; if you're looking to experiment, start with these three simple containers.

Booksonic-air (https://hub.docker.com/r/linuxserver/booksonic-air): Create your own audiobook streaming server as featured in the Holiday 2019 issue (no. 172). Booksonic-air is a new version of the server that's far more responsive, and boasts a new interface. Once running, access it through your browser at http://192.168.x.y:4040, where 192.168.x.y is your server's IP address.

Bitwardenrs (https://hub.docker.com/r/bitwardenrs/server): Host your own Bitwarden password server in a lightweight container built using the Rust programming language. You need to map port 80 to an unused port (such as 4000), and "/data" to a folder such as "/home/docker/containers/bitwarden" but that's about it for the basics. You do need to access it through a secure connection—the easiest way to do this, plus gain remote access, is by pairing it with LinuxServer's Swag (https://hub.docker.com/r/linuxserver/swag), which comprises nginx and free Lets Encrypt SSL certificates. We covered this in issue 174.

Finally, Duplicati (https://hub.docker.com/r/linuxserver/duplicati) provides you with an easy way to keep your server's precious data backed up—point the "/source" mount path to your "/home" folder to ensure all your personal data and configuration (including Docker user and files) is backed up locally and to the cloud.

Encrypt Messages With Your Pi

YOU'LL NEED THIS

PYTHON
Any sort of computer with Python installed—we used a Raspberry Pi. Plus some secrets to pass on.

SENDING secret messages has been with us for a long time. Various methods have been devised, most with a weakness. Take the very ancient idea of tattooing your message on the head of a slave or serf, as seen in antiquity. First there's the problem that slavery is morally repugnant, then you've got to shave the head, do the tattooing, and wait for the hair to grow before sending them to your co-conspirator. Then they need to be reshaved before the message can be read.

Far better to use something like, paper. There are invisible inks, of course, or there are ciphers, which change the letters used to make up the message so that only someone who knows the precise transformation can reverse it.

In 1500 BC, the Mesopotamians encrypted recipes for pottery glaze, and ciphers proliferate through ancient Egypt, Greece, and Rome. The first people to systematically document their methods, however, were the Arabs, with Al-Farahidi (710-786 or 791AD) writing the *Book of Cryptographic Messages*, considered the first book on cryptography written by a linguist.

We're bringing things a bit more up to date, with a look at how you can use—and break—the notorious Enigma cipher machine on just about any computer you've got handy. –IAN EVENDEN

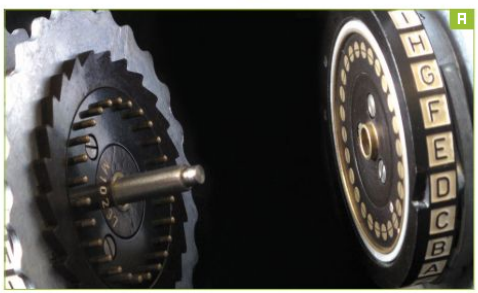

1 THE MACHINE

When not being simulated on your computer, an Enigma machine is a mixture of mechanical and electronic systems. This is 1915 technology, so there's a keyboard from a typewriter, a series of wheels through which an electrical signal can be passed, and a board of lamps that light up depending on which letter they are showing. Each rotor is different, but each has 26 electrical contacts on its edges [**Image A**], and wiring inside that connects the contacts together, swapping letters over. Stepping components ensure that at least one rotor turns with every press of a key, changing the electrical pathway for the next press.

›› The trick is that every machine shipped with a set of identical rotors. Usually, three from a choice of five were loaded into the machine, and set to their starting positions. The choice of positions was vital, as the encryption was reversible if you typed the ciphertext produced by one machine into a second machine set to the same starting positions. In World War II, the positions for each day tended to be provided to radio operators in a book, often printed on blotting paper to make its destruction easier.

›› There is more to Enigma than the rotors, though. From the depressed key (perhaps an A), the current flows through a plugboard that can be configured to swap letters over. If A is connected to D, the D signal goes into the three rotors, where it becomes perhaps a J, then a W, then a P, and is reflected back through the rotors again by a reflector, but takes a different path, becoming M, then Z, and finally F. Then it returns to the plugboard, and if F isn't connected to anything, the F lights up on the lamp board. The only rule is that nothing can ever be enciphered as itself. On the next key-press, at least one of the rotors has turned, so A still becomes D, thanks to the plugboard, but D might now become L when it reaches the first rotor, so its path through the machine is different and you don't get F at the end.

›› Cracking these starting positions, either by stealing codebooks or through remarkable feats of mental agility combined with electro-mechanical aids, was carried out by the remarkable brainiacs at Britain's Bletchley Park during the war. See pg. 46 for more on them.

2 THE SIMULATOR

These days, we have computers. The best Enigma simulator we've found is from Félix Defrance, who goes by the name omnitrogen on GitHub. It comes

THE VALUE OF CRIBS

In cryptography, a crib is a section of plaintext that is known, or can be guessed, providing a means of solving the cipher.

Many cribs come from operator error. For example, if you know that your enemy tends to end or begin its messages a certain way, this can be useful. Repeated messages reading "nothing to report" are also a good source of cribs.

If you have double agents, try getting them to send messages of your own devising to enemy HQ using the cipher—you then know both the plaintext and ciphertext, making it easier to work out that day's settings.

Weather reports provided a good source of cribs during WWII, as German enciphered reports could be compared to known conditions. The capture of the German short weather codebook (weather reports were condensed into seven-character messages that were enciphered using Enigma) during a raid on an abandoned Nazi submarine in WWII provided many cribs until it was changed.

© WIKIMEDIA

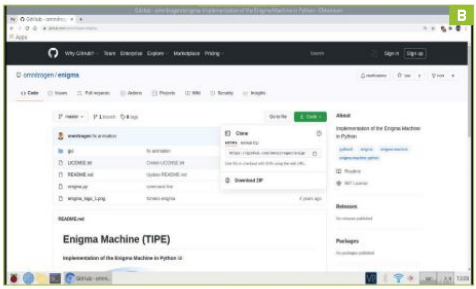

in the form of a Python script, so is relatively platform-agnostic. We've installed it on a Raspberry Pi, partly because you don't need enormous processing power, and partly because the Pi OS comes with the Thonny Python integrated development environment installed, as long as you choose the "Full" version of the OS; in Raspberry Pi Imager, look under "Raspberry Pi OS (Other)." For other platforms, you can get it from https://thonny.org. Note that, in Windows, this may trigger a warning from Windows Defender—click "More Info" followed by "Run Anyway."

3 THONNY PYTHON

Pop up the Raspberry Pi equivalent of the "Start" menu, and you'll find Thonny Python under "Programming." Don't click it yet. If it's not there, open Terminal and type `sudo apt-get install python3-thonny`. Next, go to https://github.com/omnitrogen/enigma and download the Zip file via the green "Code" button [**Image B**]. Unzip it, and delve into the "GUI" folder. Locate the file "enigma-gui.py," right-click it, and choose "Thonny Python IDE" to open it with. You should get something that looks like the image below [**Image C**]. Press "Run."

4 ENCRYPT AND DECRYPT

The Enigma machine doesn't simulate the plugboard, but gives you a selection of three rotors, and you can place them in any order in the slots, including duplicating them, using the arrows at the right of the horizontal slots. Below those is the reflector, which doesn't change. Then there are three vertical slots that represent the starting positions of the rotors. Set these how you please by dragging them with the mouse. With the rotors set 1, 2, 3, and A, A, A, the popular phrase "Maximum PC Minimum BS" becomes QSWOITCZUYYOLLEEVJ [**Image D**]. Move rotor 1 to position B, and it becomes KVVEJSTNGIPXMKMHZB. Note that there are no spaces, numbers, or punctuation used, and it throws an error if you add them. Stick to letters, perhaps using an X as a spacing character.

›› To decrypt, reset the machine to its starting settings, and feed in the cyphertext you generated. If the settings are right, you'll get readable plaintext in the output box.

5 CRACK THE CODE

To use this to communicate, you need to transmit your rotor order and starting positions to your recipient without the enemy becoming aware of them. Alternately, if, like us, you're intercepting your neighbor's Morse transmissions, you could try a brute-force attack, but it can take a long time. Breaking into Enigma is still difficult today without some sort of clue about what the message says—in 2017, an AI was used to mount a statistical attack on an Enigma message (see boxout), and took 12 minutes to solve. There is a way to do it in Python, with scripts by British intelligence agency GCHQ and a historically accurate Python library by Brian Neal, but it's very long and beyond the scope of this tutorial—find out how deep the rabbit hole goes (it calls for eight Pi machines at one point) at https://projects.raspberrypi.org/en/projects/octapi-brute-force-enigma.

USING AI

The AI that brute-forced Enigma in 12 minutes was tasked with cracking the four-rotor naval variant of the machine, code-named Shark. It used over 2,000 virtual servers—Droplets, in the terminology of provider Digital Ocean—and was trained in German using *Grimm's Fairy Tales.*

Having been taught to tell the difference between random strings of letters and real German words, the AI, called Enigma Pattern, was let loose on Enigma. As it was a brute-force attack, the AI tried every possible combination of the rotors and plugboard (15 billion), using the kind of Python recreation we've used to test its guesses at the key.

Initially running on AWS, the AI would have taken two weeks to find the answer. When it was transferred to Digital Ocean, with 1,000 Droplets, that dropped to 19 minutes at a rate of 13 million combinations tested per second. Doubling the server count to 2,000 raised that rate to 40 million per second, which is when the 12 minute time was set.

In comparison, Alan Turing's Bombe machines, which helped automate the solving of Enigma during WWII, took 20 minutes to run through all the settings for a single rotor.

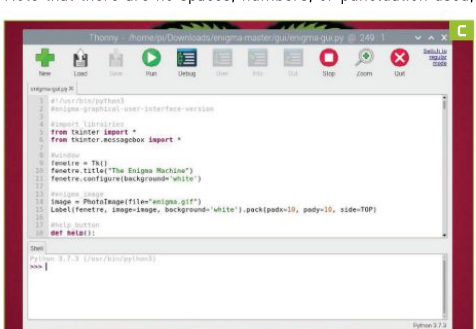

© RASPBERRY PI FOUNDATION

Emulating the Sinclair ZX Spectrum

YOU'LL NEED THIS

A COPY OF RETRO VIRTUAL MACHINE.
32 and 64-bit versions are available for Windows and Linux, plus a 64-bit package for macOS.
www.retrovirtualmachine.org

AT LAST, WE'RE FEATURING the machine we've been teasing for months: the Sinclair ZX Spectrum. To anyone outside the UK, this name may be meaningless, but in Britain the machine was a cultural icon. Far cheaper than anything Americans were exporting, the homegrown Spectrum was a working-class hero that spawned a generation of influential British coders.

Although the machine had cheap graphics and primitive sound, its rock-bottom price spawned a phenomenal budget gaming scene.

For anyone willing to look past its shoddy presentation, an anarchic game culture awaits, with classic titles unlike any other system. –JOHN KNIGHT

1 DEVELOPMENT

The ZX Spectrum was developed under the leadership of British inventor Clive Sinclair, who was renowned for making electronics affordable through ruthless simplification.

›› Originally code-named the ZX82, the Spectrum built upon the ZX80 and ZX81—machines that helped establish budget UK computing—but brought a wave of improvements.

›› Depending on the model, Spectrums had either 16 or 48KB of RAM, and were powered by Zilog's Z80A CPU. The microcomputer now had a color display; the name Spectrum was chosen to advertise its new eight-color palette.

›› Unfortunately, it still had an awful keyboard. Its cheap rubber keys were unresponsive, with an infamous "dead flesh" texture. Sound was also primitive, with just a single-channel beeper. Due to memory constraints, games were usually plagued by color clash whenever sprites crossed.

›› But heavens, it was cheap. Launching in April 1982, initial prices were £125 (16KB) or £175 for the 48KB model ($200 and $280 at 1982 exchange rates). However, prices soon dropped to the point where Spectrums were around half the cost of a Commodore 64, or a third of the price of Amstrad's CPC.

2 GAMING

While console owners had to continually save for new games, British shelves filled with £2 Spectrum titles. A thriving magazine culture featured cover tapes packed with games. Add in the system's rampant piracy, and Speccy owners soon developed large game collections. What Spectrum games lacked in style, they made up for in substance. They were often highly innovative and a goldmine for anything weird and different.

›› Although its graphics looked ghastly compared to the Amstrad or Commodore, the machine did have some tricks up its sleeve. It could outperform the Commodore in 3D games, and often had much sharper sprites compared to the blockier characters on rival systems. This sometimes resulted in richer environments, if you could look past the primitive coloring.

›› As for some gaming highlights, *Manic Miner* and *Jet Set Willy* are perhaps the most famous, and the *Dizzy* series was ubiquitous in Britain. *Exolon*, *Ant Attack*, and *Silkworm* are awesome. And for something a bit more punk, how about *Advanced Lawn Mower Simulator?*

3 HISTORICAL IMPACT

The Spectrum truly brought computing to the masses, spawning more bedroom coders than ever before. This era put Britain at the forefront of programming talent and earned Clive Sinclair a knighthood.

›› Despite strong sales, issues with other products brought financial problems. Sinclair Research sold off its computing division to Amstrad in April 1986. Amstrad continued to sell Spectrums with improvements such as better keyboards, an improved sound chip, more RAM, and built-in tape or disk drives.

›› Spectrums were also heavily cloned. Timex was officially licensed to build clones, with poor sales in the US but strong sales in Poland and Portugal. Countless other clones were produced in Eastern Europe and South America, and some clones are still being produced.

›› By the early '90s, sales were still much stronger than expected, with Amstrad focusing on 128K models as budget game machines. The Spectrum was eventually discontinued in 1992, with sales of around five million.

4 HOW DO I EMULATE IT?

There are many Spectrum emulators. We like Retro Virtual Machine (RVM), an emulator for the ZX Spectrum and Amstrad CPC.

›› Head to the "Downloads" page, where there are packages for Windows, Linux, and macOS. Windows users are given a ZIP file containing a stand-alone executable. Linux users are given a Debian package and a binary tarball, or can install RVM from the Snap store.

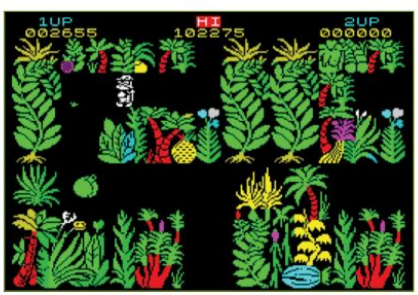

›› ***Sabre Wulf*** **(1984) nicely demonstrates typical Spectrum graphics, with the psychedelic eight-color palette, color clash, and games mostly taking place at night....**

© ULTIMATE PLAY THE GAME

›› **Non-British hipsters may recognize this diminutive machine from *Black Mirror: Bandersnatch.***

FINDING SOFTWARE

The Spectrum library is enormous and you will have no trouble finding software. Leaving alone dubious ROM sites, http://archive.org has plenty of titles, including cover tapes. Cover tapes are an incredible source of free games and often have exclusive content.

For anyone looking for new Spectrum games, a selection of modern titles are available at http://zxonline.net, http://playonretro.itch.io, and www.retroworks.es.

For an extensive back catalog, nothing beats http://worldofspectrum.org, which seems to take care in only allowing games that are either abandonware or free to distribute, and includes magazine articles, books, user manuals, and more.

5 STARTING WITH TAPES

When you first open RVM, there will be a blank interface, but don't panic. Click the hamburger menu in the top-left of the screen to open the sidebar menu. Click the "Create machine" button.

›› A menu now appears with a choice of Amstrad or Spectrum. Choose "Spectrum," and click "Next."

›› The following screen has a selection of Spectrum models. The easiest choice is the +3e model, which has both tape and disk drives attached. After selecting a Spectrum, click "Create." You will return to the main screen with your model ready to launch.

›› Clicking on the new Spectrum opens a new window with a blank emulated monitor. Press the power button located in the slide-down menu at the top of the screen.

›› Newer Spectrum models have a boot menu with the option "Loader" or "Tape Loader." Press Enter and the Loader waits for you to load a cassette.

›› Back in the slide-down menu at the top-right is a cassette icon. Press this and a large tape drive appears. Press the "Eject" button to select a tape image, then press the "Play" button to get started.

6 LOADING DISKS

To load a disk, first restart your machine with the "Machine Reset" button in the slide-down menu. Click the floppy icon in the top-right of the screen, which reveals two floppy drives.

›› You need to insert the disk in the bottom drive. Press the drive's "Eject" button, which opens a disk menu. Click "Insert," which opens a file browser. Select your disk image and click "Open." Then, with your virtual disk inserted, press Enter to select "Loader" from the Spectrum's boot menu. If the Loader program finds your disk image, it loads automatically, otherwise it prompts you to load a tape instead.

7 LOADING MEMORY DUMPS

Memory dump files (normally .Z80 files) are easy to start and load instantly. Click the hamburger menu button at the top-left of the screen, and click "Load State." Select your image file and click "Open."

›› If there is any glitching, try loading the image with an older 48K model.

8 MANUALLY LOADING TAPES

Original Spectrums don't have boot menus, so you need to enter commands in BASIC.

›› First, you need to know the name of your program. Load a tape as above, but look at the tape in the deck—on the label will be the name of the program in quotes, such as "TRAIL" for *Trailblazer*.

›› You can flick between programs with the "Fast-Forward" and "Rewind" buttons, which display different programs on the cassette label.

›› Power on your machine if you haven't already done so, and at the BASIC prompt, press J, which brings up the Load command. Enter the name of your program in quotes, but do not press Shift! The " character will appear, but the ' won't. If you have entered the command correctly, your tape will load as soon as you press "Play."

SPECIFICATIONS

SINCLAIR ZX SPECTRUM

CPU:	Zilog Z80A @ 3.5MHz
RAM:	16, 48, or 128KB
Launch price:	£125 16KB (approx. $200), £175 48KB ($280)
Graphics:	256x192, eight colors
Sound:	General Instrument AY-3-8912 (128K)
Storage:	Cassette, ZX Microdrive, 3-inch floppy
OS:	Sinclair BASIC
Released:	April 1982
Production:	1982–1992
Worldwide sales:	Five million

›› ***Micronaut One* (1987) is an example of the Spectrum's 3D advantage, still providing other-worldly experiences today.**

© WIKIPEDIA, NEXUS PRODUCTIONS, 1987

BUILD IT

ZAK STOREY, EDITOR

Virtual Renaissance

Maximum airflow, minimum form factor

LENGTH OF TIME: **2-4 HOURS** | LEVEL OF DIFFICULTY: **HARD**

THE CONCEPT

THERE'S NO QUESTION that we've long been big fans of the Phanteks Shift line-up of chassis. Over the last few years, it's certainly featured in a number of our builds. It's one of those cases that really sparks the imagination when considering the potential of a system. That eccentric monolithic design was perhaps the first case to reimagine what an ITX chassis could look like, and how it could function. It doesn't look quite so distinctive nowadays, with the likes of the NZXT H1, Corsair One, and Xbox Series X all featuring a similar sort of look and feel, but there's no denying that even today, the Shift still flaunts a certain amount of flair, even with the heavy competition it faces.

However, case designs like this do tend to suffer from one problem in particular: airflow. This is something Phanteks has addressed across the last two iterations of its Shift Air chassis, though—in fact, you can read more about that in our review this issue (pg. 82). With the announcement of the Shift 2 Air late last year, we knew we had to get one in for a build, and with our resident staffer Sam jumping into the world of the virtual this month (pg. 18), there was no better scenario in which to use it than an overkill VR rig. Building a small-form-factor PC for anyone looking to pop one in their living room or man-cave made a lot of sense.

When VR launched with the HTC Vive and Oculus Rift in 2016, system requirements were, at least for the majority of PC gamers, quite trivial compared to some of the things we'd seen prior in the realm of dedicated desktop monitors. In fact, you could get away with an AMD Radeon RX 480 for pretty much everything VR-related—a card that was yours for only $240 or so back in the day. Let's see how things have changed....

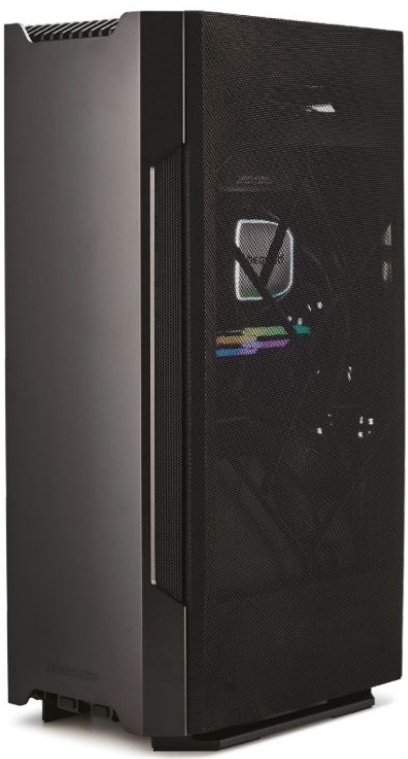

SPECS SPOT

FIVE YEARS ON, the minimum spec hasn't changed as dramatically as you might think. In fact, it remains almost the same. For the Vive Cosmos Elite, that means a Core i5-4590 or AMD FX 8350 processor (yep, that old), at least 4GB of RAM, and any form of GPU higher than a GTX 970 4GB or Radeon R9 290 4GB. For some perspective, those GPUs score on average around 8,000–9,000 in 3DMark: Fire Strike (1080p); compare that to a GeForce RTX 3060 Ti or AMD Radeon RX 6700 XT today (which should cost around $200-300 normally), which can score over three times that. But that's not where the story ends; the games themselves have their own minimum specs. In the case of *Half-Life: Alyx*, that's a Core i5-7500/Ryzen 5 1600, 12GB of DDR4, and a GTX 1060 or RX 580 with at least 6GB of VRAM. Of course, higher specs than that provide better performance, and less latency and lag.

With that in mind, we decided to throw caution to the wind, and see what we could fit in this wee new Air 2 chassis as an overkill build. We had intended to run an AMD setup (the Ryzen 7 5800X and Asus X570-I ROG Strix mobo), but the parts fell through at the last moment, and as the latest Intel parts still hadn't arrived with us, we decided to give the Core i5-10600K one last outing.

We've coupled that with a few extra choice components, including the gargantuan MSI GeForce RTX 3080 SUPRIM X GPU, 16GB of DDR4, a ton of Be Quiet! cooling goodness, and a 600W SFX-L power supply. You'll notice we're running a 1TB Samsung 980 Pro PCIe 4.0 M.2 SSD, and although PCIe 4.0 isn't supported on the Core i5-10600K, it will work at 3.0 speeds well enough until we upgrade to the Core i5-11600K and a newer motherboard.

INGREDIENTS

PART		STREET PRICE
Case	Phanteks Evolv Shift 2 Air	$100
Motherboard	Asus Z490-i ROG Strix Gaming	$230
CPU	Intel Core i5-10600K	$225
GPU	MSI GeForce RTX 3080 10GB SUPRIM X	$1,140
Memory	16GB (2x 8GB) G.Skill Trident Z NEo @ 3600 CL16	$140
PSU	600W Be Quiet! SFX L Power 80+ Gold	$195
Storage	1TB Samsung 980 Pro PCIe 4.0 M.2	$196
CPU Cooler	120mm Be Quiet! Pure Loop AIO	$85
Fans	2x 140mm Be Quiet! Silent Wings 3, 1x 120mm Be Quiet! Silent Wings 3 PWM high-speed	$76
OS	Windows 10 Home 64-bit OEM	$32
Total		$2,419

1 TEAR IT ALL DOWN

AS ALWAYS with our builds, we start things off by stripping the chassis down as far as it can go. This is actually incredibly useful from a journalistic perspective, because we can really get to grips with what's changed compared to the last generation if we do this every single time. Fancy external panels are nice and all, but it's the internal layout and frame that make or break a build. Here we have a redesigned power outlet/pass-through, better cable management options, removal of the fold-up Phanteks cable-tidy bar (now a separate addition that you can install later), better cable management at the rear of the case near the GPU and by the motherboard, plus substantially better airflow cutouts across the entire thing. All the panels are removed in a similar manner as the first Shift case: Pop open the lid at the top, and loosen the two thumbscrews holding each panel in place, then simply pull up, and each one pops off.

2 AIRFLOW, AIRFLOW, AIRFLOW

THE PHOTOGRAPHER looked at us bemused when we asked him to take a photo of this, but this is such an epic update, we couldn't let it go unnoticed. This is the rear panel of the chassis. The top (where your I/O plugs into the motherboard) is on the right, with that tiny oblong cutout, and the bottom (where your kettle lead plugs into the PSU pass-through) is situated on the left, with the hexagonal cutout. When installed on the chassis, your 120mm AIO can sit above the hexagon, and your 140mm exhaust fan near the motherboard can go further up. What's so great about this? All those cutouts. You've still got mesh running along the sides, but the bulk of the cutouts are unfiltered, meaning minimal air resistance for those fans, which significantly improves airflow. It does mean more cleaning in the long run for the end user, but performance should be much better compared to having a mesh layer sitting on top.

3 AWKWARD BACKPLATES

MOVING ON to prepping our motherboard, and first thing on the agenda is to try to sort out as much as we can before we get the bulk of the hardware inside the case. But before that, someone needs to tell Be Quiet! how to make a decent CPU cooler backplate, because this thing is nothing if not frustrating. You have to thread these screws through the back of the Intel plate, then "secure" them with a rubber o-ring on the other side. Next, with your third hand, carefully insert it through the cooler mount holes on the board, while securing it in place with the other included thumbscrews on the other side of the mobo. It isn't that complicated but, unlike most manufacturers' backplates, the screws fall out, or don't fit correctly, or decide to go for a walk if you so much as jiggle it ever so slightly. We saw better designs in 2013.

4 FINAL GOODBYE

HERE WE'VE TAKEN the liberty of installing the Core i5-10600K in its socket. To do this, lift the retention arm, raise the bracket, line the processor up with the notches/golden triangle on the socket, carefully lower it into position, lower the bracket back down, making sure to sit it under the Torx screw, then lower the retention arm back in place, and it locks into position. You can leave the plastic cover on the socket; it pops off once the bracket locks down, and means you're less likely to damage the socket. We've also removed the top heatsink off the audio componentry/M.2 installation location, ready for the 980 Pro. There are three small screws holding it in place and a tiny plastic bracket on the left to remove to do that.

7 PROBLEMATIC GPUS

NEXT UP, we installed the PCIe riser. This goes in easily, threaded through from the opposite side of the chassis. Install it as you would any other graphics card, and you'll be on track. We've yet to have a bad experience with a graphics card running off a PCIe riser like this. However, this is a very small case, and graphics cards are getting ever larger. Technically, the Shift 2 supports cards up to a maximum length of 335mm/13.19 inches, a maximum width of 150mm/5.91 inches, and a depth of 59mm/2.32 inches (think 2.9-slot card). The SUPRIM X is 336mm/13.22 inches long, 140mm/5.51 inches wide, and 61mm/2.40 inches deep.... Yeah. And we knew this going in. So, out come the screws for the entire PCIe bracket.

8 CHILLED

IGNORING WHAT'S clearly going to be a problem later, we installed pretty much everything else apart from the graphics card. First, the fans. We're running three: one for the CPU cooler, a 120mm high-RPM Silent Wings 3 from Be Quiet!, a 140mm installed here (see picture) as an exhaust, and an extra 140mm for additional intake, at the bottom. In hindsight, and due to the nature of the filtration, it may have been better to install this fan here as an intake, rather than an exhaust, as between it and the motherboard you need to install CPU power, three of your fans, plus any D-RGB components, making it easy to get a cable caught in a blade when powered on, which is less likely to happen if the rear fan guard is in the way.

5 ROOMY!

INSTALLING THE MOBO in the Shift 2 Air is fairly simple. Lay the chassis on its side, making sure any interior cables are out of the way. Place the motherboard in position and install with the screws from the Phanteks toolbox. You do need to be mindful about where your power cables are going with this case, particularly for the eight-pin CPU power. Some ITX motherboards do feature them top-left, like most modern ATX mobos, but other manufacturers have a habit of moving them around the board, which can be a problem for cable management. When it comes to ITX builds, always, always, always make sure you've got space to run your CPU power, and your motherboard and chassis support it with ease.

6 FORGETFUL!

THERE'S AN IRONY in the fact that we forgot to install the storage until now, but it's easy to do at a later date. With the RGB heatsink removed, align the M.2 drive with the notch in the socket, and the notch in the drive, then slot it in position. Our Samsung drive features a copper-laced thermal sticker to help dissipate heat, but if your drive doesn't have a thermal sticker, the M.2 SSD controller chip should face toward the heatsink. The controller acts like a mini-processor, so the cooler it is, the longer the drive can run before it throttles. Controller chips are usually taller than their NAND flash counterparts, so more of the controller makes contact with the thermal pad/heatsink; NAND chips tend to prefer hotter temps.

9 AIO WOES

FOR THOSE who've been following our adventures with the Z490-I ROG Strix Gaming, you'll remember we had trouble installing the EK-240mm AIO with it in our Razer Tomahawk build, and the same is true with the Be Quiet! Pure Loop. You can't install it in what would be considered the correct orientation, as its mounting brackets make contact with the audio componentry/M.2 PCIe mounting area and VRM heatsink. So, we rotated the mounting brackets by 90 degrees, then secured the Be Quiet! pump block on top. What's also curious about this AIO is it seems to house the pump at the bottom, next to the radiator, in its own self-contained unit. And the CPU block and the radiator pump thing have their own PWM cables that need installing. We've also installed our Silent Wings 3 120mm high-RPM fan, so it's pulling cool air in through the rad. The Pure Loop comes with a Pure Wings 2 fan, but we wanted something with a little more grunt.

10 MORE INTAKE

WITH THE AIO INSTALLED, we removed the feet panel and popped in the bottom 140mm fan. This is a neat shot as you can see where the two USB 3.1 ports are housed to the left, with their rubber grommets. This is a departure from the original Shift, which had them on the side of the tower, lacking rubber plugs. If you wanted to go minimal, the section is held on by Phillips screws, so you could remove the lot, and reduce the number of cables needed—ideal for liquid cooling or more cable-management-focused setups. With the fan in place, and secured from inside with the relevant screws, it's simply a case of popping the feet back into place, making sure to line up the correct end with the PSU cable pass-through location. This is one of the few areas of the chassis with a removable mesh filter.

11 A TIGHT SQUEEZE

THIS IS WHERE things get trickier. We've loosened the PCIe graphics card cage from the rear of the chassis, there are only two screws holding the PSU in place, and cables are everywhere. This is because we've been pushing things back and forth trying to see what fits where. The 120mm Be Quiet! AIO is pushing up against the PSU, so we had to install the cables first, then route them, then secure the PSU in position, bending the pump and its tubes away from the PSU and closer to the radiator in a 90-degree angle. We've also had a go at fitting the GPU, and came to the conclusion that the cables need pre-routing for the PCIe connectors before we even get close to installing the card in the machine. All of this could have been avoided if we'd have gone with a slimmer SFX PSU.

12 GPU TIME

IT TOOK US three test-fits to get the graphics card sorted; this shot is after test-fit number two, when we realized we needed to remove the bracket. Because the dimensions are so tight, we can only install the card if the bracket is situated at a 45-degree angle outside the case, then we can plug the power cables in, secure it to the bracket using the included screws, then carefully slide it back in. It's still a bit of a fumble to get right, but it fits. Another interesting tidbit is the fact the Pure Loop's two fan headers can be powered by an included SATA cable. One of the fan headers we believe is for the white LED in the CPU block, and the other powers the pump at a fixed RPM. Which saves us running anything directly to the board, and gives us more space to install the cables for our three fans.

13 OVERKILL

IS THIS TIGHT? Yes. Does it fit? Also yes. OK, so it's admittedly very, very, very snug, but we did manage to get the card in the build without too much of a problem. The bracket itself is secured by two screws at this point, one on the side that you can see below, and one in the rear, tying it to the center motherboard tray. The card does bulge out a little at the rightmost end here, but the mesh panel that's situated on top of it actually features a slight extrusion in the mesh anyway, before it makes contact with the panel, so it never actually touches the card, and the fans spin perfectly. Not too shabby, huh? But, yes, in hindsight, we'd definitely use a slimmer card, time permitting. We've also managed to tidy up the bulk of that cable mess fairly well here.

14 ALL WRAPPED UP

AND THERE YOU HAVE IT: one complete super-small-form-factor overkill VR/gaming build that you can pop in your man-cave or take on the road. Done and dusted. Usually, we'd show you how to install an OS or go over some software elements of the build here, but this time we've decided to pick out some highlights. The rear I/O design has always been an epic addition to the Shift series; this panel pops up at the top, and you feed all your peripherals in via the cutout to the right. There's also plenty of room to run interior cables around there. As you can just about see, we have one of the case fan cables running from one side of the graphics card to the CPU header on the mobo. If you do this, make sure you don't catch it and rip it out by accident when plugging in other peripherals.

1 **A smaller power supply is the first and biggest change we'd make here, followed up with some slicker cables.**

2 **Next on the agenda would be a better motherboard with a little more room for AIOs. That said, the orientation of the Pure Loop does look awesome.**

3 **You can actually install a stellar-looking Phanteks cable bar here if you'd like, but with the AIO tubes and all those power cables in the way, we just couldn't fit it in, unfortunately.**

4 **No surprises that MSI's biggest RTX 3080 barely fits here. The Gaming X Trio and Ventus versions of the RTX 3080 would fit perfectly, however, and give similar performance.**

THE END FOR 10TH GEN?

EVEN WITH INTEL'S LATEST launch right around the corner, we can't write off the Core i5-10600K. Although Rocket Lake is still based on the 14nm manufacturing process, availability is going to be rough early on. It's inevitable. The last year has seen an influx of people looking to get into the world of bespoke high-end computing, and with everything else going on—the crypto-boom, new consoles, new smartphones, new cars, and everything in between—newer products are going to be scarce. And here's the kicker: Right now, from a performance perspective, AMD has the advantage in the CPU market. If you want the best performance, you have to go AMD. Core counts, IPC, clock speeds, PCIe generations, even mobo designs—all of it is better on AMD and its 5000 series chips.

That said, there's very little stock out there for those parts. So, you might be the king of processors, but if you can't supply the parts, it's all for nothing. Intel on the other hand—perhaps because 14nm is easier to mass-produce, perhaps because it has a lot of stock due to everyone flooding to Ryzen—has 10th-gen parts available. And with Rocket Lake about to launch, including PCIe 4.0 support as well, the likelihood is after a week or so, stock's not going to be there either. Certainly not by the time you read this.

The thing is, though, as consumers, we're not necessarily in too tough a position on this one. At least, not on the processor front. Because we're still stuck in the middle of this processor war, we've seen significant leaps made by both camps in terms of real-world performance over the last four generations. And the Core i5-10600K in this build is still incredibly good. It has six cores, 12 threads, and seriously impressive single-core performance that will last you a good four to five years without issue. You're not going to see any throttling. It's the same on the GPU front. If you can get a card, even running it on PCIe 3.0 instead of 4.0 sees almost no difference in frame rate. We're not at a point graphically where any of these components are going to be a bottleneck. We're rendering games on Ultra at 4K, at near enough a consistent 80fps, and that's quite close to twice the resolution that our best VR headset can muster (4.9 million pixels in the Cosmos Elite versus 8.3 million pixels at 4K).

It is frustrating that these parts don't exist, but as you can see, if you can source a GPU, even with last-gen components (whether Intel or AMD), you can still put together an incredible machine that hammers frame rates all day long (we'll let the benchmark table do the talking here). OK, it might not have the fastest storage transfer times, or be the quickest at rendering 4K content in After Effects, but for the vast majority of folk, it packs more than enough punch to keep you sated for the next four or five years.

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	4,278	1,398 (-67%)
CrystalDisk QD32 Sequential Read (MB/s)	4,985	3,531 (-29%)
CrystalDisk QD32 Sequential Write (MB/s)	4,199	3,473 (-17%)
Middle Earth: Shadow of War (fps)	72	78 (8%)
Total War: Warhammer II (fps)	65	75 (15%)
Shadow of the Tomb Raider (fps)	80	85 (6%)
Red Dead Redemption II (fps)	69	78 (13%)
Assassins Creed Odyssey (fps)	54	54 (0%)
Retail Price	$3,458	$2,419 (30%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of an AMD Ryzen 9 5950X, 32GB of DDR4 @ 3600, an AMD Radeon RX 6800 XT, and a 1TB Corsair Force MP600 M.2 PCIe 4.0 SSD. All tests were performed at 4K on the highest graphical profile available.

LAB NOTES

TESTED. REVIEWED. VERDICTIZED.

PCs Make the World Better

Computing is the ultimate example of theory put into practice

ONE OF THE BEST things about computing is the way it combines theory and practice, and often leaves you with a tangible result.

Back in 2020, I wrote a tutorial in this magazine about installing a Raspberry Pi computer and an IR camera in a bird nest box, to get an intimate view of a breeding pair as they build a nest, lay eggs, and hopefully raise their young to fledglings. As I write in mid-March from my European perch, a pair of Eurasian blue tits have been seen checking the box out, and with the help of a Wi-Fi extender, the box camera is happily streaming pictures live to YouTube. Hopefully, soon everybody will be able to see my two birds on the Internet.

This is a definite result. The theory was sound, but much could have gone wrong—I went through three power supplies and a bunch of USB cables, plus a waterproof junction, before I managed to get a powerful enough electricity supply nine feet up a tree without the rain spoiling it. The thought still haunts me that I'll switch it on one morning and there'll be a fizzle as the Pi sits in a pool of water.

I thought back to other features in the magazine. The Y2K bug, derided in the press as a fraud, was very real in both theory and practice, but it was the efforts of thousands of IT engineers that made sure there was no major disaster. Then, this issue, there's the story of the invention of the first electronic, programmable computer during WWII. Now, I'm not putting myself on the same level as the Bletchley Park crew, or even the '90s engineers who headed off the Y2K problem—I'm just a writer who knows how to use a screwdriver and can tell which end of a SATA cable is which—but all through the history of computing, the same thing has happened: Theory gets put into practice, ideas become real, and the world gets a little better every time. **–IE**

Yeah, so the image quality isn't great, but hopefully we'll have more to show soon.

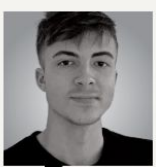

SAM LEWIS

Staff Writer

Transitioning from the world of design into technology journalism has been one heck of a positive move for me. I have always followed the industry, but recently my love for it has positively boomed. Working with and around tech has fully consumed my life for the past three months. With my partner, I've temporarily moved into an old house belonging to her family while coronavirus persists. It has been such a rewarding opportunity and the house has become a playground waiting for me to tech out. I have always wanted the space and permission to mess around with a room however I like and, finally, that has happened.

Redecorating has been the main agenda, but with this comes the prospect of combining design and technology, to create a smart home. Depending on the room, once we have redecorated it, out goes the 19-inch HD-ready TV from the mid 2000s, for instance, and in comes the obligatory 58-inch 4K television. Catch my drift?

With both my partner and myself working from home at the moment, our offices are ever-advancing, too. I am currently looking into making sure every room has a smart home element that's compatible with the Google Assistant, while also trying to keep it all on a budget.

It's been exciting to get into this position and I can't wait to see all my ideas through to make them reality.

Editor's Pick: The Wrong Program

Versatility is what really makes the PC shine

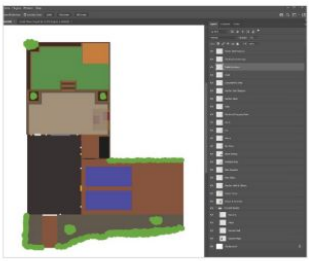

THIS EDITOR'S PICK is going to be a little unorthodox, mostly because I've wanted to talk about this for a while, but haven't found the right space to do it. Seriously, this is really going to annoy my art editor; I can see his emails arriving now!

So, I can get very obsessive over things. I've had numerous people analyze my life over the years, from family and friends to psychologists. Everything from my battles with depression and hypochondria to my passions for various sports, fitness, and things-to-do, and they've all come to similar conclusions. Basically, I pick things up quite quickly, and become very impassioned very fast, and that's partly due to having a very low level of OCD. It stems from a desire for symmetry and order, which strangely helps with how I learn things.

Now, OCD's no joking matter. It's often misportrayed in the media as being a love of cleaning, or a desire to do something perfect every single time. However, it's a condition that can make people extremely miserable. It's a compulsion, rather than a desire. Your mind tells you that you need to do it this way, otherwise something awful is going to happen. Or worse still, it makes you worry and think about things that are distressing. That in turn can manifest itself as massive bouts of anxiety or fear, which have a hugely detrimental effect on mental health. And it's all out of your control; you can't just turn it off like a switch.

Where does this desire for order go in my world of computing? Well, I live and breathe Google Sheets. I find it a fantastic way to organize things, and oddly relaxing. For instance, in my personal Drive right now, I have 33 spreadsheets. I have spreadsheets for obvious things, such as monthly finances (with three separate variants for different scenarios for the next five years). Then I have less obvious spreadsheets, like ones planning holidays or for following my favorite New Zealand rugby team, and the ridiculous extreme, where I have world-building spreadsheets for story-writing, where I'll go into detail about book plots I'd like to write one day. And that's just my personal Drive, not my work one.

Then there's my use of Photoshop. This will make anyone who's worked in publishing, CAD/CAM, or 3D design cry. Building CVs or PDFs? InDesign? Nah, Photoshop. Designing a bespoke wooden TV stand? Autodesk? Nah, Photoshop. Making an architectural plan with one pixel to one centimetre accuracy, to redevelop my garden, driveway, and home? SketchUp? Nah, Photoshop. You get the idea.

These programs are not designed to be used this way, yet they work perfectly for the job at hand for me. That's the true brilliance of the PC. For instance, take our Photoshop tutorials in the mag. Although Sam is showing you a very effective way of doing something, he's showing you his way. There's nothing wrong with that, but there's at least another 10 different ways of doing the same thing, with 10 different programs, all of which will achieve the same effect and quality. There's no one right way of doing anything. So, maybe next time you start a project that you feel you need a particular new piece of software for, or perhaps you're putting something off because you don't have that piece of tech or the skills to use it, think about taking a product or program you're already familiar with, and with a little outside-the-box thinking, I'm sure you can come up with a simpler, faster solution using it that fits your needs perfectly. **–ZS**

$Free, https://drive.google.com
$10 per month, www.adobe.com

© GOOGLE

Reviewed...

76 Intel Core i9-11900K
78 Gigabyte Aorus 17G

80 Lenovo Legion Y25-25
82 Phanteks Evolv Shift 2 Air

84 Vive Cosmos Elite

87 Corsair K70 RGB TKL
88 Endgame Gear XM1 RGB Dark Frost
89 Razer Kiyo Pro

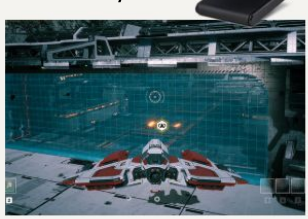

90 Everspace 2
92 Duolingo vs. Babbel

Intel Core i9-11900K

Oh boy, here we go...

THERE'S LOTS TO TALK ABOUT here, so we'll keep it brief where we can. This is Intel's latest 11th-gen Core series of desktop processors. It's also the company's seventh iteration on the 14nm transistor design. It debuted 14nm with Broadwell (fifth-gen, successor to Haswell and Devil's Canyon) in 2015, and since then the company has struggled to break past that 14nm barrier, with delay after delay proving costly. There have been several attempts, and Intel did see some success with low-powered laptop and ultrabook chips in 2019, thanks to the 10nm+ Ice Lake CPUs, but desktop parts never arrived.

As AMD has continued to push the limits of its architectural design and manufacturing processes, Intel's had to pivot to keep up, increasing core counts, voltages, IPCs, and frequencies in the process, and working its way around the engineering and design problems induced by those unorthodox methods of improving performance.

With its 5000-series CPUs, however, AMD presented something that even Intel couldn't match, with dominance in both single-core IPC and multicore capacity, and continued support for PCIe 4.0. Intel's 10th-gen series, when it launched, was a hard pill to swallow for most people wanting to upgrade, as it only represented increased core counts, with slightly higher clock speeds and IPC, more heat, and, of course, a new socket.

Something had to change, and with the multicore, IPC, and gaming crowns absent from Intel's marketing decks, a bold plan was put in place to back-port a well-established 10nm architecture to the 14nm manufacturing process. Thus Rocket Lake was born, built off the Cypress Cove architecture, itself a derivative of Sunny Cove, used in the 10nm Ice Lake laptop CPUs we mentioned.

That's led to some interesting hand-offs between the two very different platforms. Intel's still using a monolithic chip design, as opposed to AMD's modular style, but the cores, despite remaining on the 14nm process, are far denser than previous generations. Intel has also added its Xe integrated graphics (named UHD 730 or 750) to these CPUs for a gain of 40–50 percent performance dependent on model, finally included support for PCIe 4.0, with 20 lanes available, increased native DDR4 support to 3,200MHz, and added a few memory overclocking features here and there.

INTERESTING INTERFACE

More interesting are the chipset updates. With the new 500-series boards, we have a new DMI, with x8 PCIe 3.0 lanes (same bandwidth as x4 PCIe 4.0 lanes), memory overclocking enabled on H570 and B560 boards, discrete Wi-Fi 6E support as standard, Thunderbolt 4 support, and support for USB 3.2 Gen2x2. The reason for doubling the width of the DMI is likely because the boards also support 10th-gen CPUs, so it's not unreasonable to think you could install a 10th-gen chip and still get PCIe 4.0 performance from drives that aren't running directly native to the CPU, but through the DMI. Neat. But we're going to have to check that.

When we look at performance, things get even more interesting—because, sadly, it's not that impressive. In our testing, we saw figures on average about 3 percent slower than our AMD Ryzen 9 5950X in single-core applications. In multicore scenarios, well, it just wasn't pleasant to watch. Admittedly, that last part isn't entirely a fair comparison, as the 5950X has twice the number of threads as the 11900K, but all the same.... It is better compared to the previous generation of chips, by about 8-14 percent on average, but as they weren't competing anyway, well, yeah, it's a tough sell.

So, what's the positive takeaway from this? Well, it is an interesting architecture, perhaps a little lower on cores than we'd like, but it's a solid contender. Being manufactured on a 14nm process in a time where we're in a chip drought could be a massive boon to supply. As we've said multiple times, even this issue, you can have the best processors in the world, but if they're not available for sale, it's all for naught. The lower-end chips might actually be the more interesting parts—the Core i5-11400, in particular, could be ideal for a budget 1080p gaming machine. And if the architecture holds true, and Intel finally manages to deliver on its 10nm processes later this year, we could see some significant improvements, including more cores across its range. Is it going to be enough to beat AMD though? Hmmm.... **–ZAK STOREY**

BENCHMARKS

	Intel Core i9-11900K	Intel Core i9-10900K	AMD Ryzen 9 5950X
Cinebench R15 Single (Index)	257	228	**269**
Cinebench R15 Multi (Index)	2,345	2,608	**4,312**
Tech ARP's x264 (avg FPS)	50.32	54.64	**91.84**
Fryrender (m:s)	1:15	1:23	**0:49**
Total War: Warhammer II (fps)	**85**	81	84
Ghost Recon Wildlands (fps)	58	**59**	58
Price Per Thread	$34	$27	**$25**
Recommended Retail Price	$540	$530	$799

Best scores are in bold. Intel test bed consists of a Gigabyte Z490 Aorus Pro AX, 32GB of Corsair Vengeance RGB Pro DDR4 @ 3200, and an Nvidia GeForce GTX 1080. All games were tested at 1080p on the highest graphical profile, with average fps noted. AMD motherboard is the AMD ROG Crosshair VIII Dark Hero.

VERDICT **7**

Intel Core i9-11900K

+ VOLATILE Good integrated graphics; still a solid processor; PCIe 4.0 finally.

– MEEK Pricey; not enough performance; fewer cores; hot and thirsty.

$540, www.intel.com

SPECIFICATIONS

Cores/Threads	8/16
Base/Turbo Clock	3.5/4.7GHz (5.3GHz single core)
Architecture	Cypress Cove
Lithography	14nm
Memory Support	DDR4 3200
PCIe Support	20x PCIe 4.0 (x16x4)
Integrated Graphics	Intel UHD Graphics 750
TDP	125W

Lack of due process.... Get it?

Its sharp utilitarian design suggests great power beneath.

Gigabyte Aorus 17G

Strong, mean, desktop rivaling machine?

SINCE GIGABYTE debuted the world's first RTX-based laptop in 2020, it's been pushing the boundaries of performance in this form factor. With the new Aorus line-up comes high-performance GeForce RTX 3000-series laptop GPUs, and in this 17G model, we're blessed with an RTX 3070. Should be interesting.

First we have to address the elephant in the room: The 17G is huge. But is that such a bad thing? Its large size and chunkier body provide more airflow, resulting in better cooling. And the build quality is superb, with a unibody design that feels incredibly rigid and reassuringly solid, though it's a little heavy.

You wouldn't want to sling this in a backpack on your commute—even if you could fit it in; it's more of a "put it anywhere and play" size. The battery is 99Wh, lasting most of the day, but you'll need to plug it into the grid for the best gaming experience.

The 17G has been designed with competitive gaming in mind, and you can tell. The panel is a 17-inch FHD IPS LCD model, with an impressive 300Hz refresh rate. It's super-snappy and responsive, with very good color management. Sure, it's not perfect for productivity, as it's not UHD, but it's no slouch when it comes to gaming. Games feel silky smooth, ensuring the gamer has the best advantage when playing.

The graphics card and screen aren't the only standout features on offer, though. With the Aorus 17G model comes a much welcomed mechanical keyboard that features the world's first Omron B3KL switches on a laptop. The keys are very clicky, with deep travel, bringing you that desktop mechanical experience on the go. It's a joy to write with these keys, especially for long durations. The keyboard is also RGB backlit, with tons of customization options. Sitting below the keyboard is an equally solid trackpad. It responds sharply, and clicking on either side is satisfyingly tactile.

It also features a fingerprint sensor to keep your pesky peers away, and there's more security features on offer when it comes to the design of the webcam, as it has an integrated slider to cover it. Unfortunately, the placement of the webcam is not ideal, because it's under the screen, which causes your face to appear from a rather unflattering angle. We assume the reason is to create tighter bezels around the sides and top of the screen, which is an unnecessary compromise, seeing as tiny webcams have been integrated into tight bezels before.

We had to save the best till last: Performance. The RTX 3070 with Max-Q Design GDDR6 8GB is paired with the 17G's 32GB of RAM, and to top that off, it houses a 10th-gen i7 processor inside, too—anything you dare to challenge this laptop with, it deals with spectacularly. Games are exactly what this was made for; running on ultra results in over 60fps, which is the ideal goal for a good gaming experience, and the fact that the machine is portable is the icing on the cake. It's a behemoth of power.

The Gigabyte Aorus 17G is exactly what you need if you want to take a solid gaming machine with you on the go. For people transitioning from console gaming to the PC side of things, this works very well paired with a 4K TV, too, running solid detail at over 60fps—better than the likes of the Xbox One or PS4, and it still offers better performance even unplugged.

So, performance is great, the screen is top-notch, and the keyboard is excellent, yet there are a few things holding this back from perfection. The bulk, slightly outdated aesthetics, oddly positioned webcam, and only 512GB of storage, plus the fact that it isn't the most portable laptop, all work against it. But on the whole, and at this price, it will deliver a great gaming experience. **–SAM LEWIS**

VERDICT **9**

Gigabyte Aorus 17G

POWER PLAY Great gaming performance; super-responsive FHD panel; clicky integrated mechanical keyboard.

POWER STRUGGLE Not enough storage for gaming; bulky chassis; badly placed webcam.

$2,100, http://gigabyte.com

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	1,030	1,510 (47%)
CrystalDisk QD32 Sequential Read (MB/s)	3,374	3,540 (5%)
CrystalDisk QD32 Sequential Write (MB/s)	2,530	2,978 (18%)
3DMark: Fire Strike (Index)	13,610	20,015 (47%)
Rise of the Tomb Raider (fps)	92	86 (-7%)
Total War: Warhammer II (fps)	62	85 (37%)
Tom Clancy's Ghost Recon: Wildlands (fps)	49	67 (37%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our gaming laptop zero-point is the Acer Predator Triton 500, with an Intel Core i7-8750H, Nvidia GeForce RTX 2060 Max-Q, and 16GB of DDR4-2666. All game tests are performed at 1080p at the highest graphical profile.

SPECIFICATIONS

CPU	10th Gen Intel Core i7-10870H
Graphics	Intel UHD Graphics 630, Nvidia GeForce RTX 3070 with Max-Q Design GDDR6 8GB
RAM	32GB DDR4 @ 2,933MHz
Screen	17.3-inch FHD @ 300Hz
Storage	512GB SSD
Ports	1x HDMI 2.1, 1x MiniDP 1.4, 3x USB 3.2 Gen1 Type-A, 1x Thunderbolt 3 USB Type-C, 1x 3.5mm headphone jack, 1x microphone jack, 1x UHS-II SD card reader, 1x DC-in jack, 1x RJ-45.
Connectivity	RTL8125-BG Realtek (2.5G) Ethernet, Intel Wi-Fi 6 AX200, Bluetooth 5.1
Weight	5.95lb
Size	1.0 x 10.8 x 15.9 inches

It's only 1080p, but Lenovo's new 25-incher packs a 240Hz IPS punch.

Lenovo Legion Y25-25

IPS and 240Hz for a price you can afford

BIG PANEL SIZE, huge pixel count, and epic refresh rate. Good luck getting all that in a single monitor for sensible money. On a realistic budget, something's gotta give.

If you're a gamer, that something is probably resolution. 4K is nice in theory; in practice, it's a nightmare to drive all eight million pixels at remotely decent frame rates, let alone at truly high refresh rates well above 100Hz. You might also be willing to compromise on size, especially if a monitor delivers in other areas.

Enter the Lenovo Legion Y25-25—270 bucks' worth of what purports to be prime gaming panel. It's just 25 inches across, which is pretty puny by today's standards. And there are only 1,920 by 1,080 pixels on offer. But the Legion's refresh rate clocks in at 240Hz, and it's based on IPS panel technology. Now you're interested.

That combination of very high refresh rate with an IPS panel at well under $300 is certainly a head-turner. But it's not all this screen has to offer. There's also claimed 1ms pixel response, punchy maximum panel brightness of 400cd/m^2, and compatibility with AMD FreeSync and Nvidia G-Sync adaptive refresh protocols. There's also basic HDR support in the form of HDR400, but no local dimming.

The build quality is sweet, too, with a lush metal post supporting the screen, clipping neatly into the rear, and offering the full range of tilt, swivel, height, and rotate-into-portrait adjustments. The power supply is also built into the chassis rather than a separate untidy brick. Nice.

Regarding connectivity, there's HDMI and DisplayPort, plus a USB hub, and a conveniently positioned headphone jack, offering excellent cable management if you wish to attach a headset and peripherals direct to the screen.

On paper, then, all very promising, but what's the Legion like in practice? First impressions are a little disappointing. Despite the relatively high brightness rating, it's not particularly punchy. Lenovo claims 8-bit color fidelity; subjectively, the Legion Y25-25 has the feel of a cheaper 6-bit IPS panel, with its slightly dull colors and indifferent contrast.

That said, it has far better viewing angles than the obvious alternative: a TN panel. Jump in game and it's immediately obvious that the 250Hz refresh translates into very low latency. The modest 1080p resolution certainly helps. Firing up a shooter like *Call of Duty: Warzone* proves that fewer pixels makes for more haste. It feels super-speedy, even if the visuals lack the detail and sharpness of a 1440p panel, much less a 4K monitor.

Pixel response as opposed to input latency isn't so impressive. Lenovo has included three levels of user-configurable pixel overdrive in the OSD menu. The most extreme setting results in fairly obvious and nasty inverse ghosting. The middle option is the best compromise.

It's not quite as blur-free as the best TN panels, however, and doesn't deliver as sharp an experience as the 1ms claim implies. Of course, there's no industry standard for measuring pixel response, so the spec only provides a rough guide to real-world performance. Long story short, the Legion Y25-25 has decent pixel response, especially for a IPS panel, but the fastest TN panels and even some more expensive IPS monitors are quicker.

Of course, the Legion Y25-25's price is the context to those relative shortcomings. For the money, expectations need to be kept in check. This is not a $1,000 screen. For sure, a 144Hz or 160Hz monitor with a more vibrant panel will probably deliver a better overall experience for most gamers, but it's welcome to have this as an option for esports and similar, where refresh rate really is king. **–JEREMY LAIRD**

VERDICT **9**

Lenovo Legion Y25-25

BANGIN' 240Hz refresh makes for a serious snappy feel; fabulous build quality.

LEFT HANGIN' Dull IPS panel; 1080p resolution is down on detail.

$269, www.lenovo.com

SPECIFICATIONS

Panel Type	IPS
Resolution	1,920x1,080
Brightness	400cd/m^2
Contrast	1,000:1
Pixel Response	1ms
Color Coverage	99% sRGB
Refresh Rate	240Hz
VESA	100 x 100mm
Inputs	DisplayPort 1.2, HDMI 2.0

Phanteks Evolv Shift 2 Air

Super-compact madness

THE ORIGINAL EVOLV SHIFT series of chassis from Phanteks was quite a departure from regular case design. It was perhaps one of the most interesting products of 2017, and came in two variants, both of which tried to capitalize on a super-small footprint combined with a versatile design. It's no secret that we here at *Maximum PC* were fans of the two cases, featuring both of them multiple times across a variety of different builds, the last one being our liquid-cooled, pink, console-killing beastie. But that's not to say they were perfect in every way.

The biggest problem these systems had came down to a lack of available intake air and exhaust space. In fact, our console-killer's original configuration, prior to the installation of Mayhems' XTR coolant, often hit 85 C on the GPU alone. Removing the two dust filters from the side vents of the rear panel dropped that figure by a solid 10 C, and removing an entire rear glass panel of the chassis reduced that figure even further, down to 59 C, which is an utterly staggering temperature drop.

In fact, it was such a common problem that companies such as Modmymods.com still offer custom-cut Evolv Shift front panels for exactly this reason, with any number of cutouts and designs developed to help alleviate the problem—albeit at a fairly substantial cost.

Phanteks recognized this, too. At least, on higher-end gaming systems, where cooling can become an issue, it is far from desirable. To that end, in May 2019, at Computex in Taiwan, Phanteks debuted its latest variant: the Evolv Shift Air. Based on the shorter design of the two, the case's side panels swapped out the tempered glass for a mesh element instead, helping to relieve a lot of the original heat problems the chassis suffered from—but that was it. Fast-forward to the end of 2020, and we were finally graced with a new version: the Shift 2 Air.

The Shift 2 Air isn't a massive redesign of the overall layout by any stretch of the imagination, but the subtle changes Phanteks has introduced are welcome additions. The rear of the case now features an aluminum panel, complete with rectangular cutout running down its entirety. Better than that, it also features zero dust filters behind the cutouts. Dust filters typically introduce a lot of resistance, and can hinder airflow (in fact, the first thing we do here at *Maximum PC* for our personal builds is remove dust filters entirely for better performance), so their absence here is definitely a plus. Additionally, the power cable pass-through position has been tweaked ever so slightly, and the interior has seen some slight changes, too, to better help the would-be SFF builders out there. Cable management is improved, the front I/O USB 3.1 Type A ports have been moved to the bottom, and the feet bracket no longer falls off if you so much as look at it. Phanteks has also redesigned the cable tidy bar, and we found it easier to install the motherboard and its associated cables as well.

By default, it supports up to three 140mm fans; however, it'll only support a 120mm AIO at the bottom of the chassis. SFF power supply support is still the norm here, but it can get a little cramped in the bottom, so unlike our cover build this issue, we recommend you opt for a half-length SFF power supply if you can, to give you a bit more wiggle room.

When we first ordered this in for review, we assumed that, given the industry's obsession with tempered glass, the mesh side panels would look dull and drab, but with a completed build, and having spent some time with the chassis itself, it's a very clean, professional look, and it also helps to avoid any potential fingerprints. The mesh itself is a premium cloth, rather than a metallic design, and overall it's a crisp-looking case.

If you're after a super-small chassis, with a clean form factor, and an epic propensity for design and airflow, the Shift 2 Air might be the case for you. There's only a few things we can criticize, and that mostly falls down to a lack of USB Type C port, and it could perhaps do with being ever so slightly taller to allow for larger GPUs. Apart from that, though, this thing is golden. **–ZAK STOREY**

VERDICT

9

KICK ASS!

Phanteks Evolv Shift 2 Air

NITROGEN Vastly improved airflow; strong cable management; great value; aluminum design.

METHANE No USB Type C.

$100, www.phanteks.com

SPECIFICATIONS

Motherboard Support	ITX
PSU Support	SFX, SFX-L
2.5/3.5-inch Support	1x 3.5-inch, 4x 2.5-inch (2x included, one dual SSD bracket)
Max Radiator Support	1x 120mm
Fan Support	3x 140/120mm
Dimensions	7.3 x 19.3 x 10.8 inches
Graphics Card Clearance	13.2 x 5.9 inches (2.9 slots)
CPU Tower Clearance	3.3 inches

Short, compact, with plenty of airflow. Much like our editor.

Vive Cosmos Elite

Where does Vive's latest headset stand in the VR world?

HTC HAS BEEN AT the forefront of the "new-era" VR race since 2016, when it originally launched its first headset, the Vive. Since then, alongside the likes of Oculus, HP, and more, it has been constantly evolving its systems to improve the overall virtual experience. The Cosmos Elite is the result of more than five years of development time, and is the company's top game-focused headset, designed to deliver ultimate precision for elite VR gamers with SteamVR tracking support. Is this the case, though? Does it finally provide a seamless immersive experience? After all, that's the goal for virtual reality.

Setting up VR can be a pain—cables, sensors, making sure you have enough room, and so on. There's no change with the Cosmos Elite. Luckily, the numbered box design makes it easy to locate items, with the steps printed on a handy diagram. In the box, you get the headset, two controllers, two base stations, and all the cables needed to get them to work together. The whole setup, including PC and monitor, requires at least five outlets, so you might need to order another extension cord. The base stations (sensors) go on either side of the room, facing each other. The controllers are wireless, and pairing them to the headset is straightforward, but overall, setup can be lengthy, generally taking 30–60 minutes. Getting the headset fitted correctly and mapping out your play area is essential for the best experience.

We have to admire the build quality of the Vive Cosmos Elite, right down to the thick confidence-inducing cables. Every aspect of the device feels solid and hard-wearing. The controllers have a nice heft to them and are textured, with a rubberized grip, so you don't lose control mid-fight. The headset feels very durable. The base stations are glass-fronted and rubber-based, to ensure that the signal is strong out the front and they don't slide if you opt to place them on a surface instead of wall-mounting or whacking them on a tripod. When you are paying around a grand for something, this is exactly the quality you'd expect to receive, so props to HTC there.

Going back to the headset, this is one of the best the company has released to date. Comfort is a priority but, of course, it's always subjective. An abundance of padding creates a soft cushion around the head. The downside is the weight—it's not the lightest. It can cause a pull on the front of the head/face, leading to a little discomfort over time. It can also cause a bit of distraction, resulting in blurred vision. Think of it like a pair of glasses: Once suited to you, they're perfect for your eyes, but move them slightly, and your vision becomes a little blurred. The same applies here. A click wheel on the back controls the tightness, which is incredibly easy to adjust. Neat. An elasticated band on top can also be tweaked to aid comfort.

But what's all this if the experience doesn't match the high build standards? Well, gameplay is rewarding, to put it lightly. With a 90Hz refresh rate and excellent tracking, it's super-smooth, which is essential for an immersive 3D

Still sharing similar design cues as the original, some things don't need to change.

experience. The screens are dual 3.4-inch, with a resolution of 1,440x1,700, or 2,880x1,700 when combined. This is a tight PPI of 671, but it needs to be, as your eyes are so close to the display. For comparison, a 32-inch 4K monitor is around 137 ppi, but you don't have your eyes up against the screen—at least, we hope not. Once calibrated, the quality is exactly what you'd hope for. Good quality screens help the immersion and stop motion sickness, so that's a plus. Attached are two headphone-style open-back speakers that hover over your ears. They were difficult to put in place, but the sound quality was decent. A great positive for the Vive systems are the controllers; the design and feel of them make interacting in-game a pleasant experience. They have a very reactive touch wheel, and great tactile buttons, with a gripping side button. Haptic vibration responds well, and the overall immersion is superb, due to good ergonomics.

If you are looking to go fully into VR and get the most out of it, this is an impressive headset to get you started. Just remember, VR is pricey, so do your research and make sure you understand exactly what you need to power it, the price of the games, and whether you can dedicate a whole room to it. If you want excellent immersion, though, the Vive Cosmos Elite can deliver. There's just a few things holding it back, the main one being its price. **–SAM LEWIS**

VERDICT

8

Vive Cosmos Elite

+ ANOTHER DIMENSION Rewarding immersion; great quality; full VR experience.

– UNFORTUNATE REALITY Optional cost for wireless; a little heavy; very pricey; awkward audio.

$900, www.vive.com

SPECIFICATIONS

Screen	Dual 3.4-inch 1,440x1,700 (combined 2,880x1,700) @ 90Hz
Field of View	110 degrees
Sensors	G-sensor, IPD sensor, gyro sensor
Audio	Stereo headphone
Connectivity	USB 3.0, DP 1.2, proprietary connection to faceplates
OS Required	Windows 10
Weight	2.0lb

SUBSCRIBE TODAY

Find your next great read

Subscribe today and get instant access on your iPad, iPhone, or Android device. Additional print options available.

http://bit.ly/MaxPC21

Corsair K70 RGB TKL

Slick updates all around

TENKEYLESS KEYBOARDS—you'll either love them or hate them. Or, if you're like our other resident keyboard reviewer, Dave, you'll think they don't go far enough, and want to chop off the top F keys, too. This reviewer, however, falls firmly in the former camp—namely, what's wrong with a good old numpad? Well, ergonomics mainly.

TKL keyboards and their skinnier counterparts position you far better during long sessions at your desk, and it's all to do with your mouse, bizarrely. Get rid of the numpad, and you shorten the length of the keyboard, meaning you naturally bring your mouse closer toward you. In an ideal world, you should be sitting with your hands shoulder-distance apart, one resting on the WASD keys, and the other placed on the mouse, somewhere near where a numpad traditionally sits—not sprawled out across your desk. This is, we've been told, better for your posture, although there doesn't seem to be much research on this. But many a game developer has proclaimed their messed-up right shoulder as proof of this concept. And what's to blame? The numpad.

You can probably see where this reviewer's opinion lies on the matter. Yep, nothing but full-sized key-switch-laden goodness for him. But, that said, we can't ignore those who appreciate the slimmer side of keyboard life, so here we're reviewing Corsair's latest and great K70 RGB TKL, and there's a lot to love about the plucky little underdog, too. Our sample comes in a linear Cherry MX Red SKU, but it is available in MX Speed and a Silent variant as well. For those not familiar with the smoothness of an MX Red, it's very sensitive to the touch, requires minimal force to actuate, and has an actuation distance of 2.0mm (1.9mm for Silent, 1.2mm for Speed).

The K70 RGB TKL has also inherited a lot of traits from its big brother, the flagship K100 RGB, including the super-nippy AXON processor, capable of up to 8,000Hz hyper-polling, aluminum frame, PBT double-shot keycaps, and general design cues, too. And you get the full suite of dedicated media keys, volume scroll wheel, and even a detachable USB Type C cable rounding it all out. Corsair's also adjusted the font face to be a little less "gamery," and there's subtle visual cues here and there echoing the design language the company adopted in summer of last year. RGB backlighting is there, of course, on a per-key basis, but sadly still no Capellix LEDs just yet (still too expensive, we've been told, although we'll keep harassing the PRs about it).

SOLID METAL GEAR

All around, it's a solid piece of gear, with fantastic build quality, a super-slick aesthetic, and a feature set that'll make anyone happy. It's hard to criticize, and at $140, it's not too expensive either, although it could be a touch cheaper.

Another nice touch is the update to Corsair's iCUE software suite. We're running a beta version to test out the K70 TKL, and it is glorious—a massive improvement over its predecessor. We're talking a new UI, per-key assignments (both on a software and hardware level, the latter for when iCUE isn't running), complete lighting customization, and performance and device settings, all in sensible, easy-to-understand locations. Corsair's clearly spent some serious development time on making sure everything about it is easy to understand and useful, and it really does show. Dare we say it, it makes SteelSeries's Engine software look a little basic.

The K70 RGB TKL is a fantastic little board, filled with enough tech to make your eyes bleed. In its MX Red configuration (our preferred switch of the three), it's quiet, fast, and versatile enough that you'll be more than content using it from your day job all the way to late-night gaming sessions. If we had one complaint, it's the lack of Capellix LEDs, but that's just nitpicking, and would no doubt bump the price up by $40. This is an exceptional product, and with the update to iCUE, if you can afford it, you won't be disappointed. **–ZAK STOREY**

VERDICT **9** KICK ASS!

Corsair K70 RGB TKL

CONVERTED Slick design; solid build quality; super-responsive; dedicated media keys; better for posture; iCUE is excellent now.

OLD HAT Capellix LEDs, please; bit pricey.

$140, www.corsair.com

SPECIFICATIONS

Form Factor	TKL
Switch Type	Cherry MX Red, Speed, Silent
Switch Durability	50 million
Media Keys	Dedicated and scroll wheel
Connection Standard	USB 3.0 Type C to Type A
OS Software Supported	Windows 10, macOS 10.15
Weight	2.1lb
Dimensions	14.2 x 6.5 x 1.6 inches

Endgame Gear XM1 RGB Dark Frost

Raw performance and only minor niggles? We're in!

WE HAVEN'T HAD much experience with Endgame Gear or its mice, but a search reveals that it's a German company with very strong links to Caseking and sister companies. Caseking is one of the largest European PC enthusiast suppliers and, as such, has the financial backing and knowledgeable staff to create some impressive products. The Endgame Gear XM1 RGB Dark Frost seems to be one of the products that's been done right, and hopefully we'll see more of them.

Visually, the XM1 RGB looks very simple, and similar to a whole host of mice. It has the standard five-button layout, and all the switches feel crisp and satisfying. Our particular model, the Dark Frost, is semi-transparent, and once the lighting shines through, it creates some great effects. This mouse isn't built for showing off—it doesn't have any small screens, crazy number of buttons, or any gimmicks—it's simply built for performance. The materials are good quality, and if you value simplistic looks, you'll enjoy it aesthetically. The XM1 is available in a few other colors, too, but we didn't get hands-on with them.

The shape is aimed toward claw grip users, but during testing palm grip was used. For small or medium-sized hands, this is fine, but for a larger hand, the mouse is a little too low profile. It is something you can get used to, but if you have banana hands and use palm grip, this might not be as suitable for you as some other mice. Fingertip grip is fine, and the shape of the mouse doesn't take much getting used to at all.

The software is very basic indeed—you can reprogram the buttons, but there is only a very select number of additional functions and commands. You can change each button to perform a single keyboard stroke, and there is also support to add Shift, Ctrl, Alt, and Windows into the mix, but if you want anything more complicated than that, you're out of luck. We're told that there are software and firmware updates planned for the near future, so this might change imminently. It is rare that we alter a five-button mouse to anything other than basic functions anyway, but if we're comparing it to a more complete software suite, it is currently lacking.

RIPPLE EFFECT

The sensor is fantastic, however. The optical PixArt PMW3389 is one of the best, if not the best, sensors on the market at the moment, and it behaves as such. The 3389 has smoothing and acceleration by default over 1,900 cpi, but one of the options in the limited software is "Ripple Control." This enables you to turn off smoothing over 1,850 cpi, which reduces motion delay. This is a great feature, because in a raw-performance mouse, you want it to perform as fast as possible. The sensitivity has a range between 50 and 16,000, with increments of 50. This isn't too bad at all, but having recently reviewed mice with single-digit increments, it does seem like an area that could be improved if we are nit-picking. You can set four saved CPI modes, but the button to change this is on the bottom, so it can't be done on the fly.

All in all, the Endgame Gear XM1 RGB Dark Frost is one of the better mice we've used to date. Sure, it isn't a mouse that's bursting with gimmicks or RGB features, and the software is somewhat lacking, but if you are looking for a mouse that is lightweight, aesthetically pleasing, and fits the majority of grips, the XM1 could be for you. There are similar mice from other brands, but we're never going to grumble about competition. The more solid mice available on the market, the better the next series will have to be to compete. –DAVE ALCOCK

VERDICT **8**

Endgame Gear XM1 RGB Dark Frost

END GAME Great wire—almost feels wireless; fantastic sensor; aesthetically pleasing.

GAME OVER Software is basic; no macro support; CPI increases only in increments of 50; no on-the-fly CPI change.

$70, www.endgamegear.com

SPECIFICATIONS

Sensor Model	PixArt PMW3389 optical sensor
Max Sensitivity	16,000 cpi
IPS	450
Acceleration	50G
Programmable Buttons	5
Connection Standard	USB 2.0, 6 feet
Battery Life	N/A
Weight	0.18lb

Razer Kiyo Pro

Capturing the moment

WITH THE WAY the world currently is, webcams are rising in popularity. Video calls are essential to keep business on the go, to stay in touch with loved ones, or to use for the ever-growing streaming scene. The number of terrible built-in Nokia-brick-quality cams still seen in 2021 is madness. People want higher quality cameras for a more efficient and interactive video calling experience.

What Razer has delivered here is another fantastically crafted product to add to its arsenal. Build quality is vital for a brand; it allows the consumer to feel confident that their product is high quality. Unfortunately, this is becoming a dying breed, due to the harsh nature of fast, cheap, mass-produced technology. The Kiyo Pro breaks the mold of traditional webcams, however, by looking like a bridge camera lens. It's quite large, but looks great sitting on top of an ultrawide monitor.

The camera comes with a foldable L mount, a six-foot braided USB type A-to-C cable, and a lens cover. The mount hooks on to the monitor easily, with plenty of adjustability—it's covered in a soft rubber, ensuring that it doesn't slide off or scratch your screen, which is wonderful attention to detail. Once the camera is on its mount, the pivoting control is great for angling it perfectly to your needs. The L mount also enables you to place it on a flat surface if you want to use the camera elsewhere. Both the camera and the mount have a traditional tripod mounting hole if you want to detach it. The quality carries on into the cable; it's a great length and due to it being strong and braided, you don't have to worry about any fraying or breaking issues. The lens cover is another nice addition for peace of mind, or for protecting the camera when traveling.

The Pro is a plug-and-play webcam, ideal for a quick setup. However, we do advise downloading Razer Synapse, because that will enable you to get the perfect settings for your environment. The camera records uncompressed at 1080p 60fps, and features a wide-angle lens, which includes three modes: wide at 104 degrees, medium at 89 degrees, and linear at 79 degrees. Straight off the bat, the video quality looks exceptional paired with the widest angle option. It has an autofocus, but you can digitally opt for manual in the software if you'd prefer. It's a shame that there isn't a manual focus on the lens itself, because at first glance, it appears there might be. One slight issue we found was that the image looked slightly cool in its default settings, but using the Razer Synapse application remedied this quickly.

LIGHT RELIEF

To get the perfect image, you need good lighting, of course, which brings us to the original Kiyo webcam. This featured a built-in ring-light, which we would have loved to have seen in the Pro model. Luckily, the Kiyo Pro features a great HDR mode—handy if you have areas of bright light, as it brings more depth and detail to these washed-out areas. This only works at 30fps, but it's a good compromise, although it's a shame it doesn't also run at 60fps. The built-in microphone picks up audio well, but you are always going to get a better result from your headset or a stand-alone mic. It's just a good inclusion for those people who may not choose to go down those routes, appealing more to the masses.

As a webcam, it offers great all-around performance: Video quality is crisp, the autofocus works well, and the wide-angle mode gets a lot in one shot. The styling is a nice homage to more traditional camera lenses, and the build quality is superb, like most Razer products. At this price, however, a 4K sensor might be expected, along with HDR mode at 60fps, plus the ring light featured on the original Kiyo model to make it a perfect webcam. **–SAM LEWIS**

VERDICT **8**

Razer Kiyo Pro

CLEAR-THINKING Crisp, smooth video; wide-angle lens; great design; high-quality materials.

MORE TO COME Lacking 4K sensor and HDR @ 60fps; a little pricey; only got autofocus.

$199, www.razer.com

SPECIFICATIONS

Video Resolution	1080p @ 60/30/24fps, 720p @ 60fps, 480p @ 30fps, 360p @ 30fps
Image Resolution	2.1 megapixels
Still Resolution	1,920x1,080
Video Encoding	H.264 codec
Field of View	104/86/79 degrees
Focus	Auto
Connectivity	USB 3.0 A-to-C
Microphone	Stereo channel, 16-bit 48KHz audio codec
Weight	0.53lb

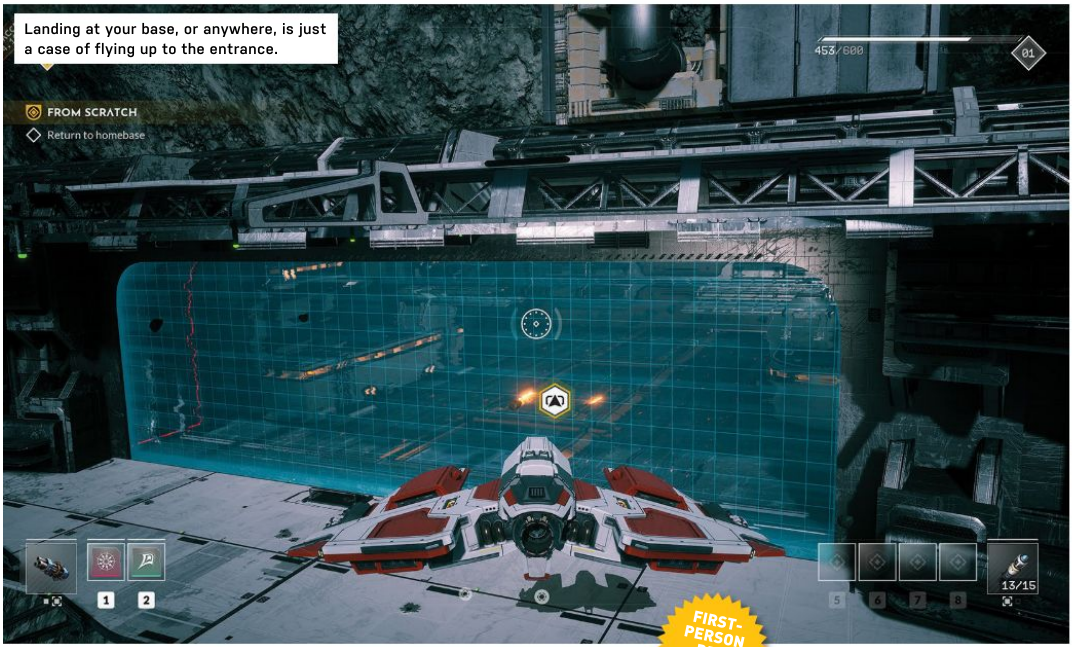

FIRST-PERSON RPG SHOOTER

Everspace 2

In space, nobody can hear you go "YEE-HAH!"

WHILE ELITE: DANGEROUS has its charms, we can't put our finger on a space shooter that's really got under our skin since *Freelancer*—and that was an embarrassingly long time ago.

And while the first *Everspace* was a fairly simple affair, in which you fought battles to accumulate credits and ship parts, but lost all except the cash if you died, this sequel has made the excellent decision to explore third-person open-world RPG territory, having Kickstartered itself into orbit.

It's also in Early Access, which means the game is not finished, and the expected completion date is sometime in 2022. Getting into a game—particularly paying money for one—at this stage is a risk, as you're likely to encounter bugs, and have your save wiped by a patch that changes something deep within the game's workings, but it also gives you the chance to experience something before anyone else, and guide its development by giving feedback.

Your character, Adam, is a clone. Produced by some sort of military program to be a natural pilot, you are the end of the line—your memories are those of all your clone-fathers, but somehow the cycle has become broken, and if you die, you die for good. That's how it's explained in one of the very early missions, anyway, but we've seen enough clones in our time to know they're an easy way of providing a player with extra lives, so we wouldn't be surprised to find a new body staggering out of a vat of goo once we've advanced far enough. That's speculation, however, as at the moment, if you die, you have to reload your last save, which tends to be back at home base before you've left for your latest mission—a system of mid-mission checkpoints would be most welcome.

Those missions are central to the game, as there's no wandering around without at least one on your list. Your base tasks you with going here, finding that, and inevitably shooting something else, while you can also pick up bounty-hunting side missions too. Character development is limited to your ship, rather than its pilot, and there's a huge range of weapons at your disposal, to be bought or found floating in space. Even as your character levels up, the perks he unlocks are reflected in your ship, not the invisible lump of meat behind the stick. The developers clearly know that what people want to see are brightly colored bolts of energy streaking toward their target, and you certainly get that, along with homing missiles, auto-cannon, space shotguns, and a degree of auto-aim that verges on the psychic (and which can be turned off in the options).

ON THE FLY

The ships are agile, capable of boosting, rolling, dodging laterally, and pitching around in looping dogfights that, we suspect, might not be particularly gravitationally accurate but look great and are an absolute blast to fly. And this turns out to be one of the great strengths of *Everspace 2*—it's simply a lot of fun to play. Ships even have *Street Fighter*-like Ultimate abilities that are charged as you destroy enemy ships, and can be released in a one-shot

© ROCKFISH GAMES

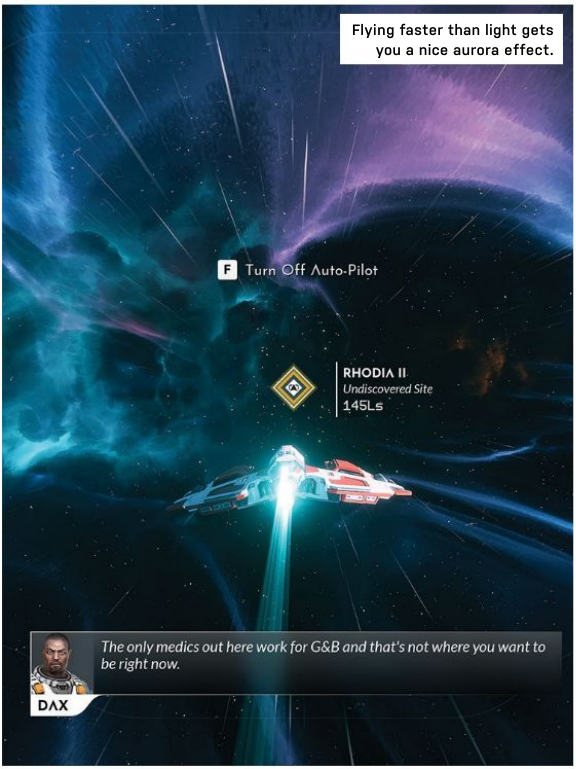
Flying faster than light gets you a nice aurora effect.
F Turn Off Auto-Pilot
RHODIA II
Undiscovered Site
145Ls
The only medics out here work for G&B and that's not where you want to be right now.
DAX

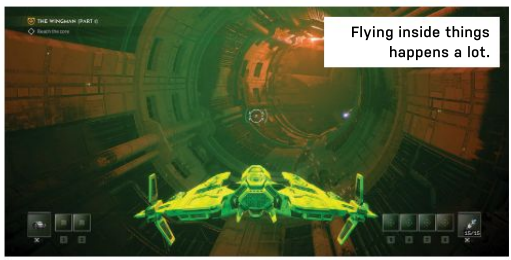

electric death supernova with the tap of a button. Abilities, which can be offensive or defensive, range from ship repairs and boosts to your power core to ship-stunning EMPs or bursts of ramming speed, and are slightly more involved: You need to press two buttons. At the same time. Don't worry, it's really not that complicated.

Lacking the simulation qualities of *Elite: Dangerous*, you will not find yourself matching rotation speeds with a spinning space station in order to dock. Everything is done with the push of a button, from docking with traders to picking up power cores and slotting them into empty sockets to complete simple puzzles that reward you with loot. Your autopilot does the job of timing your exit from FTL travel for you, meaning you should never overshoot your target as you can in *Elite*, while its in-system fast travel, known as Cruise Drive, is simply activated and deactivated with a button-press.

And while it may be heavily ship-based rather than character-focused, that doesn't stop companion characters with useful ship abilities from coming to your aid when you need them. Space itself is highly characterized and gorgeous to look at, with the crumpled wreckage of what looks like our own International Space Station spread over the orbits of giant planets, red blinking minefields like fireflies in the void, and the warping animations as capital ships' shields regenerate meaning there's always something to marvel at. Locations are hand-designed rather than procedurally generated, so while there might be a limited number of them, they're packed full of secrets. Wreckage, which is a common sight, can frequently be explored in search of loot, crumbling walls knocked down with a missile to release the goodies trapped behind them. And as a product of the Unreal Engine, performance is good, too, with the recommended system specs surprisingly mild.

What you're not going to do is build a universe-spanning empire or conquer the galaxy through trade, but the focus on combat and exploration suits us just fine. Updates to the current limited base game will be coming every few months, according to the devs. We played version 0.4.16428, but the list of additional features coming before 1.0 is long. It is easy to be positive about a game when it's unfinished, when there's no dull grind after the halfway point, or when promises that it's definitely going to improve are yet to be broken. *Everspace 2*, however, exudes a confidence and an exuberance that makes us think it's going to keep its promises. –IAN EVENDEN

VERDICT 8

Everspace 2

+ SHOULDER OF ORION Excellent-looking space shooter that's full of potential.

– SHOULDER OF MUTTON There are bugs out there; game not finished yet.

RECOMMENDED SPECS i7-4770K or Ryzen 5 1500X CPU; 16GB RAM; GTX 1060 or Radeon RX 580 GPU.

$40, www.everspace.game, ESRB: Not rated

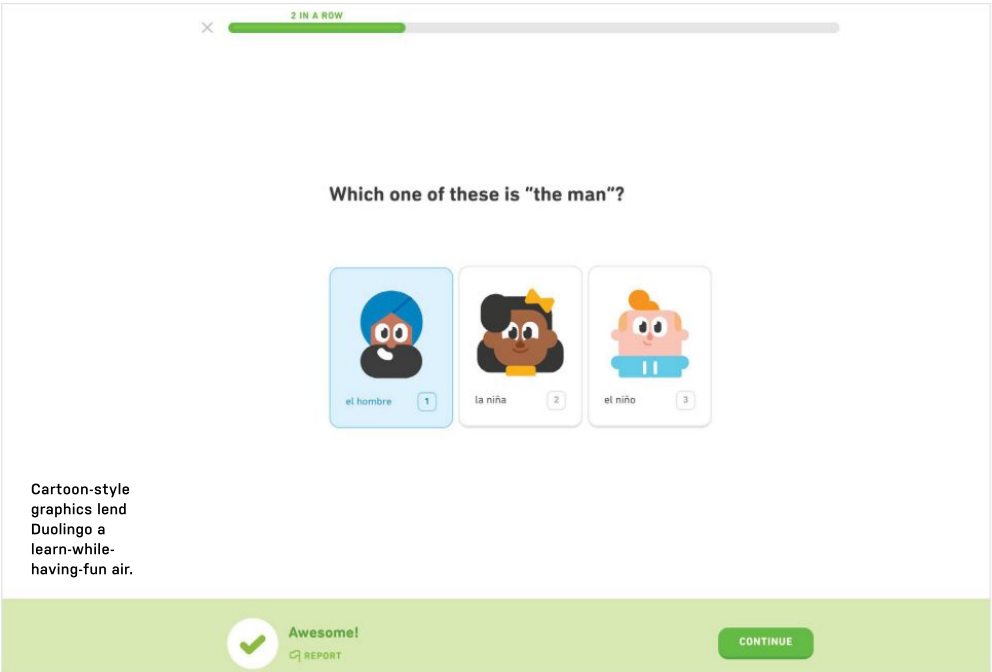

Cartoon-style graphics lend Duolingo a learn-while-having-fun air.

Duolingo vs. Babbel Learning whole new

LEARNING A NEW language as an adult can be a tricky proposition, especially if you live in a monoglot culture and don't get the chance to practice.

Apps and services that promote language-learning can come in handy here, offering daily lessons to learn new vocabulary and reinforce old lessons, so that when you do get the chance to use your new language skills, they're on the tip of your tongue.

Prominent among these services are Duolingo and Babbel. They're both popular ways of learning another language, but each takes a rather different approach.

The first major difference is in pricing: Duolingo is free. Or at least a version of it is. The free software gets you full access to its lessons—and that's it. If you want to remove ads, get offline access, make unlimited mistakes, and actually track your progress, you need to sign up for a Plus account. This costs $7 a month, and there's a two-week free trial to see whether it's right for you.

Babbel gives you a week's free trial, but makes you sign up for a subscription plan before the trial starts, so you have to remember to cancel at the end of the seven days if you decide it's not for you. There are two levels of subscription—a Babbel App sub, which gives you access to all the online streaming lessons (from $6.95 a month, if you sign up for a year), and Babbel Live, which gets you that plus access to live virtual classes (the cost depends on how many classes you attend).

The approach the two apps take toward learning a language is very different. As you can tell from the description of the subscription levels above, Babbel takes a traditional lesson-based approach, with modules and a logical structure, where basic vocabulary is taught first. Duolingo treats it as a game, with bite-sized lessons that focus on real-life goals, and leaderboards that pit you against your friends in competitive language-learning, and is organized around "skills"—short units on a single subject, such as the weather, or your family.

This is the main differentiator that will send some learners to one app and others to the other: Do you learn best in a structured environment with a "teacher" figure? Or are you more likely to retain the information in a 10-minute session on the couch with a cartoon owl? Duolingo's approach is actually a bit deeper than that. While the cartoon owl is a feature, there's a method at work behind its unusual working practices. In a study carried out by Duolingo itself, the median amount of time Duolingo learners with no prior knowledge of a language took to reach "intermediate" level was half that of college students. ("Intermediate" is defined as reaching the fifth checkpoint in Duolingo, and the A2 proficiency level for college students.)

MIND-ALTERING EFFECTS

Now, obviously, research that's been carried out by companies about themselves is always suspect, and "company says own product is good" should never be a headline, but if there's a sniff of truth that using a free app is remotely as good as doing a college course, then really we should all be doing it. Research (proper research, this time) supports a link between

© DUOLINGO

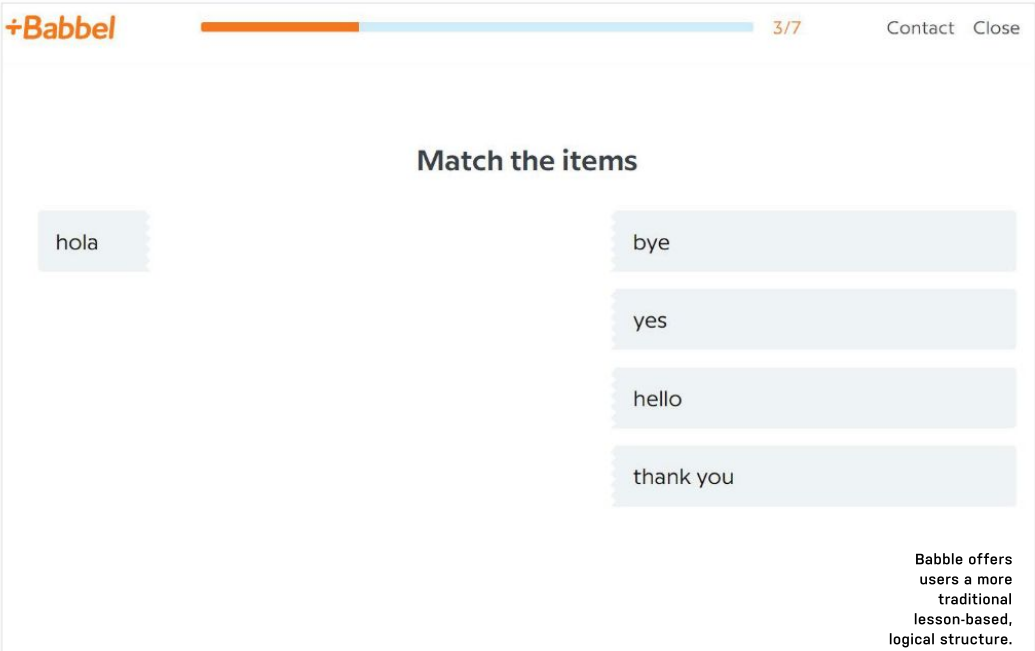

Babble offers users a more traditional lesson-based, logical structure.

languages online has never been easier

bilingualism and cognitive flexibility, plus increases in alertness, empathy, creativity, and originality. The Sapir-Whorf hypothesis suggests that the structure of a language affects its speakers' world view, and therefore people's perceptions are relative to their spoken language. So, change your language, and you literally change your mind.

That's not helping us compare these two apps, though. Babbel's interactive dialogues are based on realistic scenarios designed by linguists, with the idea that, if information is seen as relevant to the learner, it's more likely to stick in the memory. Duolingo wants you to remember what you've been taught, too, so you're frequently required to go over skills you've already mastered. Duolingo also motivates you to keep going with "streaks"—essentially, it tracks the number of days in a row you successfully hit your goals in the app—but the sentences it has you memorize are just as likely to be nonsense as they are useful, with few explanations about why it wants you to learn about snakes eating cakes, or taking your pants off—you may learn advanced words before you know the basics. Duolingo's audio, too, sounds auto-generated and unnatural, whereas Babbel's is more realistic.

Perhaps this is the reason Duolingo can offer lessons in 38 languages to Babbel's 14. The hand-crafted approach may offer more realistic scenarios, but you never know when the Spanish for snake or pants will pop into your head just when you need it. Duolingo offers instant grading for all its lessons, while Babbel uses review sessions to check that you're retaining the vocabulary.

As far as apps are concerned, Duolingo offers them for Windows, Android, and iOS, while Babbel offers them for Android and iOS only—both services can be used perfectly well through their respective web pages, however.

The differences between these language-learning services run deep, even though they help you work toward the same goal. It all comes down to how you like to learn, whether a formal teacher and classroom setup is more likely to impress the information into your mind than a videogame-like system that encourages you to chuckle at the nonsense it gets you to translate, and compete with your friends via leaderboards. Would you thrive in a calm, professional environment (Babbel), or do you prefer the instant feedback of bright colors and loud noises you get from Duolingo? That's the kind of question that free trials were made for. –IAN EVENDEN

VERDICT 7

Duolingo

+ **LINGO** Game-like structure and competitiveness enhances learning for some.

– **LINGUINE** Some sentences can be nonsense; loud approach may be off-putting.

$Free or $7 per month, http://duolingo.com

VERDICT 8

Babbel

+ **BABEL** A tailored approach to languages, with actual lessons.

– **BABBLE** Can feel a bit too much like high school at times.

$6.95 per month, http://babbel.com

© BABBEL

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> 3D Printing
> Wi-Fi Problems
> Big Rigs

3D Printing? Please!

Hi Zak, I finished reading the March 2021 *Maximum PC*. Your article about 3D printing asked if there was any interest in the subject. I'm interested!

I almost impulse-bought a 3D printer this past year, but talked myself out of it. I'm glad I did as I don't know much about it and would likely have purchased something I might regret later.

Another topic that might be of interest to the magazine is the concept of homelabs. I've purchased used systems from eBay that have 12 drive bays, dual Xeons, 10Gb NICs, and so on. I run virtualization, NAS (Unraid being my favorite), etc. I know there's been a few articles, but I would love to see more.

–Chris Rider

EDITOR, ZAK STOREY, RESPONDS: First, let me just say that since writing that article and asking for folk to email in, I have been inundated with responses regarding 3D printing, and how best to approach it. I've read every single one (although, sadly, my inbox is such a mess these days, I can't respond to everyone, although I am trying!), and we've now scheduled some content concerning those plucky plastic printers.

It was a toss-up between two of our writers—Dave Alcock and Kris Butterill—and Kris won out, so next issue, we've got a fantastic feature on getting started with 3D printing and how it all fits together in 2021. Dave hasn't been left out in the cold, though, as he decided to build his own bespoke keyboard from the ground up using a few aftermarket parts and his own 3D printed designs, so he's also writing a feature for us on that. I'm really excited about both. Obviously, I'm aware that this is *Maximum PC*, not *Maximum 3D*, but that said, we'll likely run a tutorial series over the next few issues with some cool print ideas and setup guides.

I've actually been blown away by what you can do with the medium. My 3D printing passions have expanded a bit since I last spoke about it, and I'm now buying professionally designed 3D models (to paint for my Warhammer armies) from Etsy.com and other online sources. The quality is impeccable.

Regarding homelabs, we've toyed with the idea for some time, but it's difficult to pursue—our usual PR contacts don't typically have sample availability, and we lack the budget to just buy stuff of that caliber. That said, I've wanted to do a server-rack build for some time, so it's definitely something we're looking at. It sort of echoes our Data Scientist build way back when. Perhaps when Christian's back, I'll get him to look into it properly from a hardware point of view, and I'm sure Nick Peers will have some things to say regarding the virtualization side of things. We'll keep it in mind.

Asus Router Problems

Hi Zak, I am a *Maximum PC* subscriber and I read your article regarding your dropout issues with your Asus RT-AX88U router.

My main router is the Asus RT-AC88U. I'm not sure if the AX is a newer version, but I experienced all sorts of dropouts, wireless disconnects, and so on as well, even when the router was replaced with a new one from Asus.

The problem is that this model router gets extremely hot and burns out components over time. I've had mine replaced under warranty at least three times over a two-year period by Asus, and I finally solved the problem myself by purchasing a double cooling fan (http://bit.ly/RouterFan) and setting the router on top of it.

I tried powering the fan through USB ports and this was a bit unreliable, and if the power is lost, the fan turns off, so now I'm powering it through a hefty USB phone charger that is rated at least 1A, and have not had the inconsistencies, dropouts, lockups, and hardware failures, and I no longer need to reboot the router daily.

Asus never admits that the problem is the heat generated from this router, it just replaced it whenever it started to fail.

Like you, my home network is critical, so I also purchased a reconditioned RT-AC88U to save as a backup, as I have smaller Asus routers in a wireless mesh, and connectivity to

↘ submit your questions to: comments@maximumpc.com

the Internet is crucial to WFH. If I had to replace, I'd use my ready backup, import the configuration file, and away I'd go.

–Michael Carter

EDITOR, ZAK STOREY, RESPONDS: That's interesting. I never considered that it could be heat-related. I updated to a newer firmware after writing that article, and it seemed to resolve all the problems, but it was winter. Since then, I've had two more updates, the most recent ones in February and March, and again I'm experiencing dropouts—sporadically, and only when streaming content from my LG TV from certain providers. It definitely was worse in summer, and I suppose there are processors inside, so it would make sense that they'd throttle or overheat. Fortunately, it hasn't burned out yet.

Thing is, I'm reluctant to try swapping to another manufacturer because the setup process is such a nightmare. Not so much on the router side of things, but every connected device in the house would need reconnecting to the new router, and it's a real pain. It's amazing—we have all these apps and software to transfer data from one device to another, one phone to another, and yet swapping to a new router is like dragging yourself over hot coals while clutching a bag of puppies. Take my Nest outdoor camera. It's mounted 10 feet off the ground, and requires you to press a tiny reset button to connect it to a new router. Thing is, you have to take it off the wall, and spiders like living in the mounting mechanism (I don't have arachnophobia, but I hate spiders with a passion).

That said, it might be a necessary move this summer if things continue like this. I'm just holding out for a decent router with 10GbE ports and 6E support as standard, so I don't have to upgrade again for some time. Netgear's Nighthawk RAXE500 has my attention right now (though it only has 2.5GbE), but something a little more aesthetically pleasing could sway me.

Mega Machine

Hi Zak, I thought I'd send a follow-up. I'm the guy who wrote about a computer case to hold my drink! Over the new year holiday, I built my new PC and wanted to share some thoughts from a first-time builder.

I went with a Fractal Design XL case—it continues to serve as a drink holder. I like the black monolith look. I live in a studio apartment, so don't need a Christmas-tree lighting effect, plus this case addresses heat concerns, and it came with sound dampening. I mounted my four hard drives at the bottom, next to the PSU. The case came with three fans—two in front and one in back; I added one more in front. I recommend Fractal—I broke one of the plastic clips for the side panels, and Fractal sent me two replacements. The case is a little over 22 inches high, and very modular.

Tip: Remove the back fans first and check to see if the motherboard has stand-offs installed. In my case, they were already installed and I made the noobie mistake of installing on top of them. The mobo was askew in terms of blackplate alignment and I tried to slide it below the fan but the damage was done (or I got a bad board)—the sound components got damaged: no output, or hissing static. Reinstalling audio drivers didn't help, so I ordered a replacement. This time—using the existing stand-offs, removing the back fan first—I installed the new motherboard without issue. It's a Gigabyte Z390 Aorus Ultra—my thanks to Christian G. as he detailed the install of a similar board a few weeks back. I used his steps to add my first M2 drive—a Samsung Evo 970. Can't beat the ease of no additional cables to install on an M2 drive. One thing to mention: When the PC starts up, there is a noticeable coil whine. I kept thinking it was my fans or the graphics card. Finally, I read it's because of the high-end VRMs Gigabyte uses on higher-end boards. Luckily, it goes away after a couple of minutes.

Jerry's Fractal XL build—clean as a whistle!

Continuing the theme of fewer lights, I originally had a Cooler Master RGB build and sent it back as I couldn't figure out the myriad cables and useless RGB. I ended up ordering a Fractal Design Celsius Plus S28 280mm AIO cooler. All cables in the pump, and only one cable to install on board—couldn't have been simpler. Speccy shows temperatures in the 30s C, so I guess I'm OK. And there's minimal lighting.

I used Arctic MX4 thermal paste, and made the mistake of applying too much, but it didn't fry my board. I applied a small X when I changed my board.

I installed 64GB of Corsair memory, but don't see any difference from 32GB. Have to figure out why Speccy is showing memory as unknown brand and memory speed as 1333.

I'm continuing to use my MSI 1060 6GB card as there is no hope of buying a 3060 Ti or 3070 at a reasonable price at present.

This was a huge learning experience—building my first computer after 25 years of upgrading prebuilt computers—and I wanted to thank you and your team.

–Jerry

EDITOR, ZAK STOREY, RESPONDS: Glad you got a decent system sorted, Jerry, it looks awesome, and seems like a solid rig. Coil whine can be annoying. I have problems with GPU coil whine in certain games but I can ignore it.

I agree about lighting. Don't get me wrong, you can build a classy LED-illuminated rig if you like, but it's easy to end up with a rainbow mess. RGB cabling is the bane of my existence. When I was building Dream Machine in 2019, the case had 11 RGB fans, and each batch of three required its own RGB command module, all of which needed to be daisy-chained together and connected to the mobo. It took three hours to do the cable management, and it still looked a mess.

THE BUILDS

ANOTHER MONTH of drought, another month of ridiculous prices. Unfortunately, there's no longer any stock of graphics cards at their recommended retail prices, so what we're going to do is just list what GPU we would recommend you buy if stocks were available at their RRP, and tell you what that price should be, then compare it to what you would actually have to pay if you went out and purchased it from the scalpers.

That does mean things are a little different from usual, and it's more difficult to compare these builds against those from last issue, but going forward, it should give us a nice base to work around. Sadly, this drought doesn't look likely to abate any time soon, and with the fabs that are being built taking years before they start producing chips, we could be in this position for a long, long time.

That said, there is some good news. Well, sort of. In the AMD build, the 450W Corsair CV power supply has come down in price by around $20, and we've made some savings on the storage, too. Sadly, the processor has gone up by $5, and the price of our OEM OS keys have gone up by $7, too. We've also had to swap out the memory and motherboard for the cheapest options we could find, in a vain attempt to keep the prices as low as we could.

On the Intel side, things haven't fared much better. Again, we've ditched any semblance of quality componentry with regard to the mobo and memory, and just opted for the cheapest we could get our hands on. In worse news, the Core i3-10100 has gone up by $40 compared to last issue, no doubt in apprehension of the latest Rocket Lake launches. It could be a good sign—that Intel's shifted its 14nm manufacturing over to the new architecture—or conversely, third parties could already be snapping up these parts as a precaution in case the latest 11th-gen chips are in short supply as well. The worst part about the situation is that the current Intel release schedule doesn't list any 11th-gen Core i3s yet either.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Corsair 4000D Airflow	$95	
PSU	450W Corsair CV 450 80+ Bronze	$29	
Mobo	MSI B450M PRO-M2 MAX NEW	$65	
CPU	AMD Athlon 3000G	$106	
GPU	AMD Radeon RX 6700 XT 12GB NO STOCK	$480	$1,289
RAM	16GB (2x 8GB) Crucial Ballistix @ 3,200MHz NEW	$72	
SSD	250GB Samsung 980 Pro M.2 PCIe 4.0	$80	
HDD	1TB Seagate Barracuda 1TB 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	
Approximate Price:		**$1,004 or $1,813**	

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Corsair 4000D Airflow	$95	
PSU	450W EVGA 450 BR 80+ Bronze	$50	
Mobo	ASRock B460M-HDV NEW	$94	
CPU	Intel Core i3-10100	$157	
GPU	Nvidia GeForce RTX 3060 12GB NO STOCK	$330	$1,039
RAM	16GB (2x 8GB) G.Skill Aegis @ 3,200MHz NEW	$76	
SSD	250GB Samsung 970 Evo Plus M.2 PCIe 3.0 NEW	$60	
HDD	1TB Seagate Barracuda 1TB 7200 NEW	$45	
OS	Windows 10 Home 64-bit OEM	$32	
Approximate Price:		**$939 or $1,648**	

Sponsored by newegg.com
TAKE IT FROM A GEEK.™

AGAIN, WE'RE RECOMMENDING the graphics card we believe is best suited to the build's position in our charts (aka 1080p, 1440p, or 4K gaming at 60fps plus), and giving two prices: RRP and street price.

GPU situation aside, we've still not got that many positives to talk about this time around. Our trusty Ryzen 5 3600 has seen price increases, and now sits at $220. We've not heard anything official from AMD yet, but it's telling that even these older 12nm parts are now suffering from the chip shortages, too, as folk work their way down the historic product stack in a desperate attempt to buy a "new" processor. We would recommend the Ryzen 5 5600 in its stead, but the newer chip, although available, currently comes in at $100 more than the 3600, and we just don't think it's quite worth spending the extra cash, especially when the 3600 comes with plenty of PCIe 4.0 lanes anyway. Performance is better, sure, but is it better value? Not really. We've also managed to grab a slightly more affordable motherboard ($5 cheaper, in fact), courtesy of an updated variant of ASRock's Phantom Gaming 4. There's not a huge difference between the two—in a board-to-board comparison, the PCB layout is almost identical—but the heatsinks have shrunk, and you lose the DisplayPort in the rear I/O. The cooler's ever so slightly cheaper, too, so that's in our favor.

On the Intel side of the table, things are a little more chipper. The Core i7-10700K has come down in price ever so slightly, and we've managed to shave some more cash off here and there, too, thanks to a Seasonic 80+ Gold modular power supply and, amazingly, a new, cheaper Z590 motherboard from ASRock. Same model name as the one from last issue, different chipset, and with it comes an expanded DMI bus, Wi-Fi 6E as standard, and better support, too. We've also plumped for some slightly spicier memory, at 3,600MHz, but if you're not using this system for anything more than gaming, we recommend dropping to a 2,400-3,000MHz kit.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Lian Li PC011-Dynamic	$168	
PSU	650W Corsair RMx (2018) 80+ Gold **NEW**	$100	
Mobo	ASRock X570 Phantom Gaming 4S **NEW**	$140	
CPU	AMD Ryzen 5 3600	$220	
Cooler	240mm Cooler Master Masterliquid ML240L RGB V2	$75	
GPU	Nvidia GeForce RTX 3070 8GB **NO STOCK**	$500	$1,600
RAM	16GB (2x 8GB) G.Skill Ripjaws V DDR4 @ 3,600MHz **NEW**	$80	
SSD	500GB Sabrent Rocket 4.0 M.2 PCIe 4.0	$90	
HDD	1TB Seagate Barracuda 1TB 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	
Approximate Price:		**$1,450 or $2,550**	

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Lian Li PC011-Dynamic	$168	
PSU	650W Seasonic FOCUS Plus Gold 80+ Gold **NEW**	$95	
Mobo	ASRock Z590 Pro4 **NEW**	$160	
CPU	Intel Core i7-10700K	$321	
Cooler	Cooler Master Masterliquid Lite 240	$55	
GPU	AMD Radeon RX 6800 XT 16GB **NO STOCK**	$650	$1,800
RAM	16GB (2x 8GB) Corsair Vengeance LPX @ 3,600MHz **NEW**	$90	
SSD	500GB WD Black SN750 M.2 PCIe 3.0	$69	
HDD	1TB Western Digital Caviar Blue 7200	$45	
OS	Windows 10 Home 64-bit OEM	$32	
Approximate Price:		**$1,685 or $2,835**	

RIDICULOUS PRICES ASIDE, even if everything was fine, and there were plenty of chips to go around, we still wouldn't recommend the RTX 3090. It's a beast, but the performance gain over the 3080 is 10-15 percent for a card that costs near twice as much, and has hefty flaws regarding its backplate and memory cooling. Aftermarket solutions from AIB partners have improved on that, and if you intend to water-cool it, EKWB provides options for the Founders Editions, but still, that thing runs hot.

Elsewhere, the builds haven't changed a great deal. We're seeing slight price increases in memory, storage and power supplies; it's not a good time to be a tech enthusiast or needing an upgrade. Amazingly (and we really don't understand why), the Core i9-10900K hasn't increased in price, but rather decreased slightly. Given its multicore prowess, and extra cores versus its newest sibling, the Core i9-11900K, we recommend picking up the 10th-gen variant over its new big bro if you're focused on any form of multithreaded workload that doesn't leverage Intel's Xe graphics processing tech.

We've also taken the opportunity with the Intel build to grab one of the newer Z590 motherboards. Not only will this give you the opportunity to upgrade to an 11th-gen processor if you desire, but it also comes with Wi-Fi 6E support, and if you read our feature on that last issue, you'll know it's going to be seriously important in years to come. Buyers beware, however: We haven't had confirmation that Intel's 12th generation of chips will be backward compatible with these boards. If (and this is a big if, because we've been saying it for five years), Intel manages to crack 10nm, and deliver on its promises this year of finally launching 10nm desktop chips with a derivative of the Cypress Cove architecture, there is no guarantee that they will operate in these boards. You'd hope so, but rumors are abound of DDR5 and PCIe 5.0 support—only time will tell.

As enthusiasts right now, our best advice is to tinker with what you have. Think about your cooling solutions, look at your peripherals. Sometimes you gotta make do with what you have out in the bush.

AMD INGREDIENTS

PART		PRICE	STREET PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$147	
PSU	750W Phanteks AMP 750 W 80+ Gold NEW	$110	
Mobo	Gigabyte X570 Gaming X ATX	$190	
CPU	AMD Ryzen 7 5800X	$449	
Cooler	Corsair iCUE H150i Elite Capellix 360mm AIO	$190	
GPU	Nvidia GeForce RTX 3080 10GB NO STOCK	$700	$2,300
RAM	32GB (2x 16GB) Team T-FORCE DARK Za DDR4 @ 3,600MHz NEW	$159	
SSD	1TB Sabrent Rocket 4.0 M.2 PCIe 4.0	$150	
HDD	6TB Western Digital Blue	$130	
OS	Windows 10 Home 64-bit OEM	$32	

Approximate Price: $2,257 or $3,857

INTEL INGREDIENTS

PART		PRICE	STREET PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$147	
PSU	750W Seasonic Focus Plus Gold 80+ Gold	$125	
Mobo	Gigabyte Z590 GAMING X NEW	$210	
CPU	Intel Core i9-10900K	$470	
Cooler	NZXT Kraken X73 360mm	$180	
GPU	AMD Radeon RX 6900 XT 16GB NEW	$999	$2,300
RAM	32GB (2x 16GB) G.Skill Trident Z Neo RGB DDR4 @ 3,600MHz NEW	$200	
SSD	1TB Corsair MP400 M.2 PCIe 3.0	$115	
HDD	6TB Western Digital Blue	$130	
OS	Windows 10 Home 64-bit OEM	$32	

Approximate Price: $2,608 or $3,909

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US, Inc., 11 West 42nd Street, 15th Floor, New York, NY 10036, USA. Website: www.futureus.com. Periodicals postage paid at New York, NY, and at additional mailing offices. Newsstand distribution is handled by Curtis Circulation Company. Basic subscription rates: one year (13 issues) US: $27; Canada: US$45; Foreign: US$45. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R128220688). PMA #40612608. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to *Maximum PC*, PO Box 2024, Langhorne, PA 19047, USA. Future US, Inc. also publishes *MacLife*, and *PC Gamer*. Entire contents copyright 2021, Future US, Inc. All rights reserved.

Reproduction in whole or in part is prohibited. Future US, Inc. is not affiliated with the companies or products covered in *Maximum PC*. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of *Maximum PC*. Products named in the pages of *Maximum PC* are trademarks of their respective companies. PRODUCED IN THE UNITED STATES OF AMERICA.

PAPER POWER

67% of the energy demand at U.S. pulp, paper and paper-based packaging mills is met using renewable, carbon-neutral biomass energy.

Discover the story of paper
www.lovepaperna.org

LOVE PAPER

Source: American Forest & Paper Association, 2018

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and "World's Most Advanced PCs" are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.